NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

THESIS

INVESTIGATING NAVY OFFICER RETENTION USING DATA FARMING

by

Aurel N. DeHollan

September 2015

Thesis Advisor:	Samuel E. Buttrey
Second Reader:	Thomas W. Lucas

Approved for public release; distribution is unlimited

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			*Form Approved OMB No. 0704–0188*
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC 20503.			
1. AGENCY USE ONLY ***(Leave blank)***	**2. REPORT DATE** September 2015	**3. REPORT TYPE AND DATES COVERED** Master's thesis	
4. TITLE AND SUBTITLE INVESTIGATING NAVY OFFICER RETENTION USING DATA FARMING			**5. FUNDING NUMBERS**
6. AUTHOR(S) DeHollan, Aurel N.			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000			**8. PERFORMING ORGANIZATION REPORT NUMBER**
9. SPONSORING /MONITORING AGENCY NAME(S) AND ADDRESS(ES) N/A			**10. SPONSORING / MONITORING AGENCY REPORT NUMBER**
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. government. IRB Protocol number ____N/A____.			
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited			**12b. DISTRIBUTION CODE**

13. ABSTRACT (maximum 200 words)

The allocation of nearly 30% of the Navy's budget to personnel costs, and the importance of manning fleet requirements to maintain operational readiness create a critical need for the Navy to effectively manage the size of the force. The Navy's personnel planners use the Officer Strategic Analysis Model (OSAM) to project officer end-strength based on policies, plans, and historical loss rates. The application of data farming to this model allows for investigation of different scenarios that can provide insight into both the behavior of the model and the behavior of the officer corps under various conditions. This study uses Design of Experiments (DOE) techniques to develop and implement an experimental design that determines the degree of stochastic variation in OSAM and explores the effect of a three-year period of poor retention of Unrestricted Line (URL) officers in paygrades O3 through O6. Analysis of results across multiple replications of a single design point indicate that OSAM produces very little stochastic variation. Regression modeling of the results allows planners to accurately and precisely predict the effect of this poor retention scenario on specific groups. This predictive capability provides the opportunity for proactive approaches to solving potential retention problems.

14. SUBJECT TERMS manpower, data farming, design of experiments, retention, simulation			**15. NUMBER OF PAGES** 85
			16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT Unclassified	**18. SECURITY CLASSIFICATION OF THIS PAGE** Unclassified	**19. SECURITY CLASSIFICATION OF ABSTRACT** Unclassified	**20. LIMITATION OF ABSTRACT** UU

NSN 7540–01-280-5500

Standard Form 298 (Rev. 2–89)
Prescribed by ANSI Std. 239–18

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release; distribution is unlimited

INVESTIGATING NAVY OFFICER RETENTION USING DATA FARMING

Aurel N. DeHollan
Lieutenant, United States Navy
B.A., University of California, Berkeley, 2003

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN OPERATIONS RESEARCH

from the

NAVAL POSTGRADUATE SCHOOL
September 2015

Approved by: Samuel E. Buttrey
Thesis Advisor

Thomas W. Lucas
Second Reader

Patricia A. Jacobs
Chair, Department of Operations Research

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

The allocation of nearly 30% of the Navy’s budget to personnel costs, and the importance of manning fleet requirements to maintain operational readiness create a critical need for the Navy to effectively manage the size of the force. The Navy’s personnel planners use the Officer Strategic Analysis Model (OSAM) to project officer end-strength based on policies, plans, and historical loss rates. The application of data farming to this model allows for investigation of different scenarios that can provide insight into both the behavior of the model and the behavior of the officer corps under various conditions. This study uses Design of Experiments (DOE) techniques to develop and implement an experimental design that determines the degree of stochastic variation in OSAM and explores the effect of a three-year period of poor retention of Unrestricted Line (URL) officers in paygrades O3 through O6. Analysis of results across multiple replications of a single design point indicates that OSAM produces very little stochastic variation. Regression modeling of the results allows planners to accurately and precisely predict the effect of this poor retention scenario on specific groups. This predictive capability provides the opportunity for proactive approaches to solving potential retention problems.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF ACRONYMS AND ABBREVIATIONS

AIK	Aikaike Information Criterion
BOY	Beginning-of-Year
CNP	Chief of Naval Personnel
DOD	Department of Defense
DOE	Design of Experiments
DOPMA	Defense Officer Personnel Management Act
ES	Enlisted Supply (Model)
GUI	Graphical User Interface
HR	Human Resources
NFM	NOPPS Forecasting Model
NOLH	Nearly Orthogonal Latin Hypercube
NOPPS	Navy Officer Personnel Planning System
NPS	Naval Postgraduate School
OPNAV	Chief of Naval Operations
OSAM	Officer Strategic Analysis Model
POCR	Probationary Officer Continuation and Redesignation
Q-Q	Quantile-Quantile
RCMOP	Requirements-Drive Cost-Based Manpower Optimization
SAG	Strategic Actions Group
SEED	Simulation Experiments & Efficient Design
URL	Unrestricted Line
VBA	Visual Basic for Applications
VV&A	Verification, Validation & Accreditation

THIS PAGE INTENTIONALLY LEFT BLANK

EXECUTIVE SUMMARY

The Navy's readiness and fiscal health depend on precise management of the size of the force. Required by law to keep end strength within congressionally mandated numbers, the Navy's manpower planners rely on models and algorithms to make predictions of end-strength. The Officer Strategic Analysis Model (OSAM) makes end-strength projections for the Navy's officer corps, incorporating historical loss rates and various guidelines and policies. This study explores the use of OSAM to investigate a potential retention crisis among Unrestricted Line (URL) officers. Applying data farming techniques to OSAM, we simulate a set of poor retention scenarios to gain insight into how the model performs, and how the URL community responds. The results demonstrate methods to expand OSAM's use and build projection models to help the Navy confront future retention challenges.

Running under Microsoft Access, OSAM allows the user to project officer inventory for one to seven years, by setting parameters for accessions, losses, transfers between communities, and promotions of officers. It is a time-stepped agent-based simulation that takes as input the inventory of officers at the start of the fiscal year, and according to the user-set parameters, produces end-strength values for all officer groups for the end of each fiscal year in the simulation. Losses are typically programmed to follow historical loss rates from a period of the user's choosing, with additional capability to introduce forced losses for specific groups of officers. Although the model has capacity for stochastic variation, current practice relies on single "deterministic" runs.

Data farming refers to the process of creating a data set representative of a large-scale problem by applying Design of Experiments (DOE) methods and using high-performance computing. The wide range of possible scenarios that a poor retention scenario could entail for various officers groups make it infeasible to study the problem in a comprehensive fashion using the existing OSAM platform. With support from the Simulation Experiments & Efficient Designs

(SEED) Center for Data Farming at the Naval Postgraduate School (NPS), we create programs to automate OSAM, giving manpower analysts the ability to run complex experimental designs. The experiment in this study simulates a three-year period of poor retention by introducing forced losses for three consecutive years for the four largest URL designators in paygrades O3 to O6. End-strengths for an additional three years beyond the poor retention period are also included in the data. The Nearly Orthogonal Latin Hypercube (NOLH) design used for the experiment varies the losses in each officer category to efficiently achieve maximum coverage of the possible range of values. Each design point receives ten replications to reveal any stochastic variation in the model.

The results of the simulation runs show that although OSAM in its current form demonstrates very little stochastic variation, changing the inputs in a systematic fashion can yield a wide range of results. The end-strength of the URL communities examined in this study vary significantly in the years following the poor retention period, but do so in a consistent manner that reveals certain trends and patterns. Although all officer groups suffer significant losses, end-strength for officers in paygrade O3 tend to rebound quickly once the additional forced losses stop. End-strengths for officers in paygrades O4 through O6, however, fail to increase towards their original sizes.

The complexity of the results makes it difficult to broadly characterize the relationship between different categories of forced losses and end-strength. Once the system is analyzed in terms of specific groups of officers, however, linear models allow planners to accurately and precisely determine the effect of various poor retention scenarios. The validity of the models relies on assumptions regarding the timing and scope of the poor retention period. These assumptions may not hold true in all cases, making further exploration of different scenarios an important area for follow-up research. Using the software tools produced in this study, planners now have the ability to design experiments that can examine these possibilities.

ACKNOWLEDGMENTS

I owe a debt of gratitude to several individuals who assisted me in writing this thesis. Foremost, I would like to thank my advisor, Prof. Sam Buttrey, and my second reader, Prof. Tom Lucas, for their guidance, advice, and patience during this process. I am grateful for all the assistance given to me through the SEED Center, including Professor Susan Sanchez for providing material support, Mr. Steve Upton for providing his constant assistance to run my simulations, and Ms. Mary McDonald for assisting with the analysis. I would also like to thank the folks at OPNAV N1Z4 for their hospitality during my visit and continued support, in particular LCDR Will Corley and Mr. Paul Soutter.

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION

A. BACKGROUND

Given the significant and growing proportion of Department of Defense (DOD) resources consumed by personnel costs and the direct effect on readiness of unmanned requirements, accurate and precise knowledge of end-strength is critical to the U.S. Navy's manpower management. For Fiscal Year 2014, the Navy had an estimated end-strength of 323,600 (including 53,400 officers), at a cost of $45.4 billion out of a total budget of $155.8 billion (Department of the Navy, 2013). End-strength predictions are continuously updated throughout the course of the fiscal year and play a central role in the Navy's manpower planning. Planners rely on these projections to achieve their objectives in the areas of budgeting, resource allocation, operational readiness, and occupational community management.

By law, the Navy's total number of active-duty personnel at the end of the fiscal year must conform to the end-strength-guidelines set by Congress. The allowable margin of error is small. The Navy must remain within 3 percent above and .5 percent below authorized end-strength (Department of the Navy, 2015). A highly developed set of policies allows the Navy to regulate personnel numbers closely; however, external factors and individual behaviors insert uncertainty into the process.

Within the staff of the Chief of Naval Personnel (CNP) at Navy headquarters (OPNAV N1), the Strategic Actions Group (SAG) uses the Officer Strategic Analysis Model (OSAM) as a modeling tool to develop end-strength forecasts for the Navy's officer corps. OSAM is a time-stepped agent-based simulation that takes as inputs the inventory of officers at the start of the fiscal year, and based on the inputs from various personnel plans, projects the end-strength for the end of the fiscal year and beyond for up to seven years. The results provide a snapshot of officer end-strength by designator, paygrade, years of commissioned service, and total years of service. OSAM–generated reports incorporate mandated end-strength constraints in the form of Officer Programmed Authorizations (OPA) that represent desired targets for projected inventory. Due to OSAM's comprehensive coverage of officer inventory and ability to fine-tune adjustments to simulation inputs, planners may use OSAM to simulate a variety of scenarios as a means of investigating the effects of proposed policy changes. Although the model has the capacity for stochastic variation, current practice relies on single deterministic runs.

B. PROBLEM STATEMENT

Beyond its contributions to end-strength planning, OSAM can aid in the development of policies that will help the Navy deal with the challenges of a period of low retention of its officers. Given the uncertainty in future economic conditions and unpredictable aspects of human behavior, the Navy remains vulnerable to scenarios in which unexpectedly large numbers of officers decide to leave the service. Recent surveys and studies have explored some current issues affecting retention and indicated a potential for acute problems (Snodgrass & Kohlmann, 2014). The CNP has publicly brought the issue into the spotlight and highlighted it as an organizational priority (Moran, 2014). Questions remain on what a period of poor retention might look like in terms of both immediate and long-term effects on officer inventory. Due to the diversity of the Navy's officer corps, the complexity of the issue can quickly grow beyond the scope of existing analytical methods. With dozens of discrete groupings of designator, paygrade, years of service, and additional attributes, the officer corps contains numerous potential avenues of investigation with respect to losses and retention. OSAM's complexity presents a challenge in exploring large-scale problems and thus provides an opportunity for using Design of Experiment (DOE) techniques to make tackling retention issues feasible and more efficient. For a topic such as retention that tends to focus on qualitative factors, OSAM and data farming offer analytical tools to explore the possibilities in a rigorous fashion. Data farming incorporates "simulation modeling, high-performance computing, experimental design, and analysis to examine questions of interest with large possibility spaces," making retention analysis through end-strength projection an appropriate application (Horne & Meyer, 2010, p. 1).

C. THESIS PURPOSE

This thesis seeks to provide a quantitative approach to understanding the effect of poor retention on the Navy's officer corps by using data farming to run a large set of simulations representing the widest possible range of selected input parameters. By developing a robust and efficient experiment design, we gain an understanding of the limits of OSAM's performance. From the results, we gain insight into the limits of officer retention behaviors. Metamodeling provides a framework for interpreting those results. Analysis of the results of a broad survey of scenarios can indicate what factors have the greatest impact in determining the end-strength of critical communities.

D. RESEARCH QUESTIONS

Applying the capabilities and framework of OSAM to the retention issue, the scope of an experimental study can encompass any number of specific officer communities. In this study, Unrestricted Line (URL) officers constitute the group of interest, further limited to officers in paygrades O3 through O6. This group provides the advantages of (1) forming one competitive category for promotion purposes, thus establishing relationships between end-strength for different designators, (2) providing a large data set in which patterns may more readily be identified, and (3) having an explicit connection to readiness as URL officers represent the Navy's organizational front line.

Ultimately, the intent of this study is to examine how the officer corps responds to a period of low retention. Within this context, the following questions guide the experiment design and analysis of the harvested data:

1. Under what conditions do officer communities experience retention problems?
2. What is the impact of a sustained period of poor retention?
3. How much does the response vary?

Insights gained from answering these research questions can inform the strategies and policies that the Navy may implement in response to a retention crisis.

E. METHODOLOGY

This thesis applies data farming techniques to OSAM and uses SEED (Simulation Experiment & Efficient Designs) Center high-performance computers to run the designed experiment. The SEED Center for Data Farming is an organization within the Naval Postgraduate School that promotes research and advancement of simulation analysis, particularly for defense applications (https://harvest.nps.edu). Initial preparations focus on identifying the factors within the model that have the greatest potential for affecting the response. With URL officer end-strength as the response for this study, numbers of forced losses, representing additional attrition, were chosen as the factors to be varied in the experiment. The choice of factors helps guide the decision on the type of design. The number of factors and their ranges of values, along with consideration for the simulation run-time and available computing resources, determine the design dimensions. The implementation of the experiment requires adapting OSAM to run on a computing cluster, including translation of the design into input plans for initializing the simulation. Output of the model results requires further modification to allow consolidation of the results from multiple runs. The analytical portion of this study dovetails from data farming techniques into statistical methods. After exploring relationships between factors and developing an overview of the response surface, metamodeling is used to build a comprehensive understanding of the retention issue and answer the research questions.

F. BENEFITS OF RESEARCH

Through both preparation and execution of the experiments, this study can assist Navy manpower planning by enhancing the capabilities of OSAM and providing quantitative data to add to policy discussions regarding officer retention. The development of the OSAMFarmer and OSAMRunner tools used to produce and run designed experiments give planners the ability to run more comprehensive sets of scenarios, even without access to cluster computing. Through the experiments conducted in this study and potential future research, OSAM users will have the ability to test and evaluate the model's capabilities and limitations.

G. LITERATURE REVIEW

Proper management of manpower resources is a primary concern of any organization and has prompted development of numerous models. Typically, the scope of these models goes beyond end-strength projection and into resource optimization. Although the end-strength numbers are themselves essential planning inputs, the ultimate goal is to fully and efficiently allocate personnel to manpower requirements. An overview of military manpower models reveals a wide range of purposes, designs, and uses.

1. Department of Defense Manpower Models

Modeling such a complicated phenomenon as human behaviors presents many challenges even in the context of military personnel systems with its extensive set of rules that provide a strong measure of control over these behaviors. The lack of understanding of relationships between variables and observed attributes has resulted in varying approaches (Schank, Harrell, & Thie, 1997). Consequently, a diverse assortment of models can be found across DOD agencies, varying in purpose, scope, methodology, and scale. A common feature in many of these models, however, is the primary importance of strength management. Whether for models designed to track retention, recruiting, training, or skill utilization, end-strength is usually a key input, and the results of these models often influence the design and implementation of models focused more purely on inventory projection (Schank, Harrell, & Thie, 1997).

2. Modeling Methods

OSAM's design reflects decisions that do not necessarily represent a consensus or unified convention for manpower models. A fundamental split of strength management model types occurs between aggregate and disaggregate models. Aggregate models simply calculate the total end-strength without regard to attributes of the individuals flowing through the system. The model essentially boils down to the equation:

EndStrength (or FutureStrength) = CurrentStrength + Gains – Losses

Disaggregate models, on the other hand, simulate the movement of personnel based on individual characteristics such as occupational specialty and length of service. Conceptually, the basic equation used for aggregate models governs the simulation for disaggregate models as well, but the calculations are made iteratively through the inventory of personnel groups. These models quickly become complex and thus difficult to implement with limited computing resources, but provide a much higher resolution picture of the state of the force (Schank, Harrell, & Thie, 1997). Since its entities are defined by multiple attributes, most significantly designator and paygrade, OSAM is an example of a disaggregated model.

DOD manpower models further differ in the mechanisms behind their simulations with differences representing conscious design decisions. Modelers have choices with respect to several key aspects of the design based on the following factors identified by Schank et al.:

> **a. Dynamic versus Steady-State.** OSAM is an example of a dynamic model since it projects end-strength from year to year, using the results from the previous year as inputs for the projection of the next year, which may have its own set of conditions. A steady-state model would assume that the conditions for the simulation remain the same for each year and needs to only

produce a single set of results for the end-point of the simulation period.

b. Group versus Entity. Group models place individuals with the same attributes in a single category ensuring the model treats them identically. Entity models, such as OSAM, maintain the distinctions between individuals and treat them separately.

c. Deterministic versus Stochastic. Deterministic models have no random effects and produce the same output for a specific input. Stochastic models incorporate randomness and can produce varying distributions of results for the same input. Although OSAM has the capacity for stochastic simulation, current practice hard-codes the random seeds within the model, making it a de-facto deterministic model. Measures such as confidence intervals on a mean provide an indication of how much an outcome may vary. Furthermore, deterministic models could be misleading in some cases. Deterministic models most likely do not give the same average outcome as an equivalent stochastic model (Lucas, 2000).

d. Historic versus Econometrically Adjusted. Loss rates have significant impact on end-strength, so accurate end-strength predictions depend greatly on an accurate determination of future loss rates. On the premise that past behaviors are reasonably accurate predictors of future behaviors, historical loss rates may be used for future year projections without regard to particular events, economic factors, or policies that may have affected past rates. Econometrically adjusted rates incorporate detailed analysis of economic factors, such as unemployment and the effects they would have on personnel deciding whether or not to leave the service. OSAM uses historical loss rates; however, econometrically adjusted rates developed through a separate process outside the model could easily be incorporated. (Schank, Harrell, & Thie, 1997, pp. 20–23).

The list of design choices above is not exhaustive and additional variations for both large- and small-scale adjustments could lead to an even greater divergence of models. Although two models may have the same purpose, differences in design choices could be responsible for differences in results. Understanding OSAM's design in the context of these choices may help explain patterns observed in the results.

3. Optimization Applications

Inventory projection models often serve as a prerequisite to further applications in manpower planning. Organizations not only want to know the numbers of personnel they will retain, but also what skill sets these personnel will have and how the workforce would compare to labor requirements. For the Navy, the inventory must not only stay within end-strength limits, but must also provide a balanced supply for fleet manpower requirements. Shortfalls in inventory can negatively affect readiness while overages can create budget problems. These types of constraints provide a suitable opportunity to apply optimization techniques. Clark's (2009) study used linear optimization to develop the Requirements-Driven Cost-Based Manpower Optimization (RCMOP) model that projects future officer inventory, with the objective of minimizing unmet requirements, subject to budget constraints. Optimization models offer useful guidance in developing goals for future plans, but changes in behaviors could detract from the projections' validity. Although OSAM does not provide a roadmap for planning force structures, its projections are grounded in reality and therefore provide a robust guideline for making force-planning decisions.

4. Data Farming Applications

Published literature contains a great volume of information on the design and implementation of military manpower models, but discussion on the application of data farming to these models remains limited. Previous Operations Research theses by Sibley (2012) and Erdman (2010) provide two examples of applying DOE methodology to end-strength projection models.

This thesis continues the research begun by Sibley in the application of data farming to OSAM. Sibley's study examined whether loss rates for officers who laterally transfer differed significantly and sought to determine a reasonable range of loss rates for accurate projections (Sibley, 2012). Using a previous version of OSAM, Sibley designed an experiment that varied lateral transfer rates and loss rates via adjustment factors built into the model. The scope of the study was limited to two officer communities: Surface Warfare Officers (SWO) and Human Resources (HR) Officers. The study's methodology could apply to a different or larger set of communities. The flexibility in OSAM's design allowed Sibley to introduce a new designator representing SWOs prevented from laterally transferring to another community. These officers tend to leave the service at higher rate; thus, incorporating higher loss rates in the simulation may improve accuracy. Even though the experiment included less than 10 percent of the total officer inventory, the design still required 90 factors: loss rates for each paygrade (O1–O6) for each of the three designators, and for each of the five projection years. Data farming enabled an experiment design that could efficiently and feasibly explore such a large-dimensioned space. The results of the study demonstrated the robustness of the model. Sibley concluded that varying loss rates separately for denied lateral transfer applicants did not significantly change the projections. Experimental results did suggest, however, that varying loss rates year to year in the simulation results in more accurate projections (Sibley, 2012).

A U.S. Army inventory projection model, the Enlisted Specialty (ES) model, has also been the subject of data farming application. The ES model takes into account the paygrades, occupational specialties, and years of service of enlisted soldiers, and the authorized positions available. Using historical data, the model projects future inventory of the enlisted force and produces an optimized distribution that seeks to minimize the difference between inventory and authorization. Erdman's application of data farming focused on fine-tuning the optimization aspect. The design varied the objective function coefficients in the optimization model to ensure the minimization did in fact produce the optimal result (Erdman, 2010).

Concurrent with this study, an additional thesis applies data farming to OSAM to examine the effects of changes to the Navy's Probationary Officer Continuation and Redesignation (POCR) policy (Borozny, 2015). Instead of varying forced losses to simulate poor retention as done in this study, Borozny's thesis varies numbers of junior officers released from active duty under POCR authority, and other factors, to investigate potential solutions to the Navy's over-execution of end-strength authorizations (Borozny, 2015).

5. Retention Models

Military retention models typically focus on econometric and demographic factors. Using historical data sets and surveys, researchers use regression models to identify predictive factors for retaining personnel. The majority of retention models cover only enlisted service members since the re-enlistment decision point provides a useful mechanism for tracking the effect of economics and policies on behavior (Weiss et al., 2002). Thus, inventory projection models themselves have not played a major role in retention research. In this respect, this study offers a proof-of-concept approach to try to model retention behaviors by applying data farming to an inventory projection model. Although OSAM does not have the advantage of having empirical data to support its results, the model is well grounded in real-world data with historical loss rates, actual inventories, and accurate representation of policies within the simulation.

II. MODEL

This chapter provides additional background on OSAM, including a description of its development, architecture, and explanation of the model's design and methodology. OSAM is a self-contained application for design, running, and analysis of simulations. For this study, the experimental design and results analysis portions use external programs, and while all simulations run exclusively within OSAM, modifications are made to make the program compatible with cluster computing.

A. DEVELOPMENT

OSAM was developed in 2007 by LMI, a government-consulting group, for N1 SAG's predecessor, N14 (Manpower, Personnel, Training and Education Catalog, 2015). Although OSAM has not formally undergone the DOD's Verification, Validation, and Accreditation process (VV&A), development of the model included accuracy testing. The model's continued use and upgrades confirm its value to officer strength management along with other simulation tools that are used by planners and officer community managers.

The original version of OSAM was written as a Microsoft Visual FoxPro executable application. Policies that limit access to software used on government computers resulted in OSAM residing on a stand-alone computer. Inputs to the simulation resided in 60 separate database files that were edited individually, making modification of these inputs a tedious process. Additionally, the user bore responsibility for documenting the changes. Preparation time for a single simulation typically ran 15 to 20 minutes (Sibley, 2012).

B. SOFTWARE ARCHITECTURE

The latest version of OSAM runs on Microsoft Access. Contractors from SAG Corporation translated the code into Visual Basic for Applications (VBA), bringing several benefits. Chiefly, OSAM may now run on government computers, greatly expanding access to Navy manpower planners and thus encouraging its use and further development. The separate databases now have been merged a single database with separate tables organized in a relational structure. Changes to inputs and data no longer require direct editing of the tables for most types of simulations.

C. MODEL DESIGN

As an inventory projection model, OSAM follows the fundamental equation of end-strength being equal to beginning strength minus losses plus gains. The flow of this calculation is summarized in Figure 1. OSAM takes the beginning inventory and during the simulation, according to the model settings, determines how many officers to subtract based on losses, how many offices to add based on accessions, and how many officers to promote and transfer. The output at the end of the simulation is the end-of-fiscal-year inventory. If the simulation is set to continue, the end-of-fiscal-year inventory becomes the starting point for the next time step. Since OSAM is a disaggregate model, this process applies to all the communities within the officer corps. The multidimensional characteristics of each officer entity means OSAM must also keep track of years of service, promotions, and lateral transfers. Longevity increases to an officer’s years of commissioned service and total years of service occur automatically. Promotions and lateral transfers occur according to specific plans determined by user input. Losses, accessions, promotions, and transfers comprise the entirety of events that affect the size and makeup of the inventory.

Figure 1. Diagram of OSAM model flow.

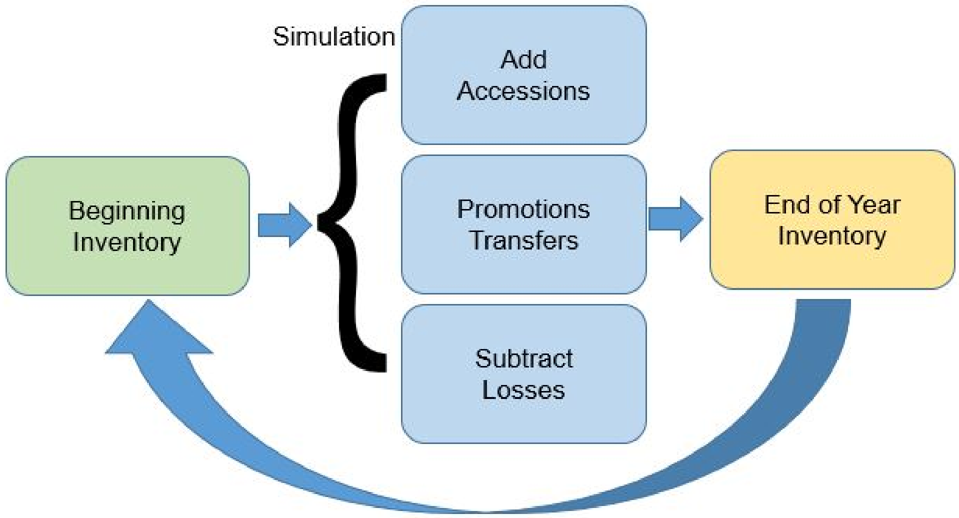

D. REGULATORY CONSTRAINTS

Another dimension of complexity to the model arises out of the numerous laws, rules, and policies that govern Navy officer management. The algorithms that determine the movement of personnel through the simulation must follow these guidelines. The Defense Officer Personnel Management Act of 1980 (DOPMA) sets the primary controls on how many officers may promote and when, and what level of opportunity they should have to advance. Additional laws found in Chapters 33 and 33A of Title 10 of the U.S. Code provide guidance on appointment and separation of officers. (Defense Officer Personnel Management Act of 1980, 2012).

Within the DOPMA and Title 10 framework, the Navy internally develops policies to manage officer end-strength and ensure fleet manpower requirements are met. The division of the officer corps into different designators allows for career specialization. Community management, including the lateral transfer of officers between designators, ensures that the officer corps remains healthy and balanced, and that each community is optimally manned. The accession, promotion, and transfer plans that serve as inputs to OSAM must follow these rules and regulations. Losses occur for both policy and natural reasons, which results in OSAM using a different approach for handling losses than other inputs. OSAM separates these types of losses to maintain a clearer picture of how different losses ultimately affect end-strength.

E. MODEL ALGORITHM

The algorithms that govern OSAM are best understood as in terms of sets. Although the data reside in a multiple tables within a database, OSAM processes the inputs by building sets of data. This section provides a detailed description of the portions of the model relevant to the experiment, borrowing from OSAM's technical documentation (Mundy, 2014). More in-depth coverage of other aspects of the model can be found in the documentation as well.

1. Losses

The treatment of losses is the most complicated aspect of the model and it has undergone significant revision since the original version. The current method for calculating losses represents the third major change to the procedure. Under this method, losses fall into three categories: (1) Natural, (2) Force-Outs, and (3) User-Added.

a. Natural Losses

Natural losses include all losses that occur due the application of historical loss rates to the beginning year inventory. They are represented by the equation:

$$\text{Losses}_{\text{proj}}(g,d) = \text{LossRate}(g,d) \times \text{Inv}_{\text{BOY}}(g,d)$$

where $\text{Losses}_{\text{proj}}(g,d)$ is the projected number of losses in paygrade *g* and designator *d*, $\text{LossRate}(g,d)$ is the applied loss rate, and $\text{Inv}_{\text{BOY}}(g,d)$ is the inventory at the beginning of the projection year. Natural losses make use of historical loss rates. OSAM calculates these rates in a straightforward manner by dividing historic losses in a year by the beginning inventory from that year. This calculation is repeated for all paygrades, designators, and years of service, giving each officer grouping its own rate. Although OSAM is not an econometric model in the sense of the user having the ability to directly adjust the model settings to account for economic conditions, users can accomplish this purpose by changing the historic loss counts used to calculate rates. Separate econometric models, such as the Navy Officer Personnel Planning System (NOPPS) Forecasting Model (NFM) can provide estimates of loss counts that feed into the OSAM loss rate calculation (Mundy, 2014).

When a user desires to alter loss rates, the changes are typically not made directly to the historical loss rates, but rather are made by applying loss adjustment factors. The default setting for a simulation has all loss adjustment factors set to one, meaning no changes to historical rates. OSAM applies any changes by multiplying historical loss rates by the adjustment factors. Factors greater than one increase losses and factors less than one decrease losses.

b. Force-Outs

Force-Outs refer to losses that occur due to Navy policies such as failure of selection for promotion and mandatory retirement due to age. OSAM implements Force-Out losses by determining a pool of eligible officers, and applying a force-out factor to remove a certain proportion from the inventory. Eligible officers are identified using a conditional statement:

$$\left[g \geq 4 \text{ AND YCS} \geq \text{YCS}_{\text{RETIRE}}\left(g,d\right) \right] \text{OR} \left[g \leq 3 \text{ AND YCS} \geq \text{YCS}_{\text{MAX}}\left(g,d\right) \right]$$

The conditional statement reflects current policy that forces officers to retire if they are in paygrades O4 and above and have reached the high-year tenure mark, or if they are in a lower paygrade and have reached the maximum allowable years of service (Mundy, 2014).

c. User-Added

The User-Added losses category consists of the additional losses programmed by the user for specific officer groups using OSAM's force-shaping input tool. The user selects which fiscal year, designator or community, grade, and year group will receive additional losses in addition to losses incurred due to historical losses and Force-Outs. OSAM allocates the losses using a pseudo-random process.

F. INPUT SETTINGS

The user runs simulations in OSAM by selecting desired settings for the adjustable inputs. These inputs correspond to the categories of changes that affect officer inventory. OSAM's scenario guides the user through the process.

1. Interface

The Graphical User Interfaces (GUI) implemented in the latest version of OSAM, version 3.1.4, represent another key upgrade that has improved its usability. Although the code and databases reside in Access, the user typically does not need to edit these components directly. A menu, labeled as Scenario Editor and displayed in Figure 2, guides the user through the building, editing, saving, and running of a simulation.

With the Scenario Editor, the user builds and edits the parameters of the simulation. Setting of the inputs occurs via selection of various plans from the Parameters tab, or selection of specific methodologies from the individual input category tabs. Once the settings are finalized, the scenario is saved and the program creates an initialization file containing the scenario parameters.

Figure 2. Screenshot of OSAM Scenario Editor.

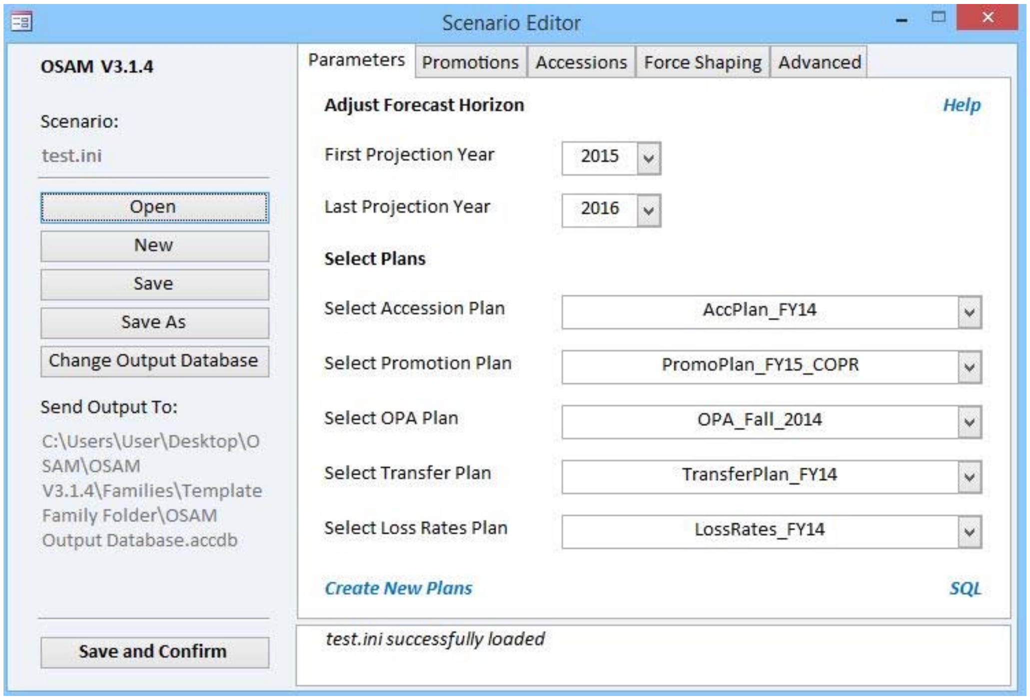

2. Inputs

From the basic input categories of any inventory projection model, OSAM has expanded the selection to account for the detailed tracking of different officer groups, and the ability to change the scenario settings to reflect various policy changes. Figure 3 shows an overview of the model inputs with the primary categories in the center and the specific setting choices on the left side.

Figure 3. Diagram of inputs into OSAM.

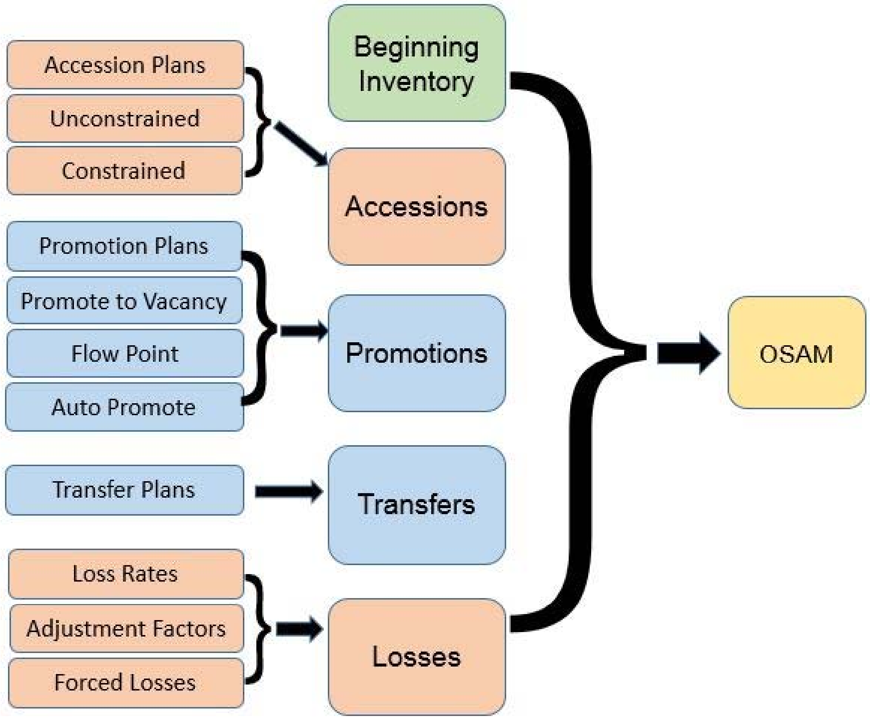

a. Beginning Inventory

OSAM contains the current inventory of active duty commissioned Naval officers with the exception of Chief Warrant Officers. It includes officers in all designators and paygrades O1 to O6. Each of these officers represents an individual entity within the simulation, but no Personally Identifiable Information (PII) is contained in the data set. The beginning inventory is updated by refreshing the Access databases and cannot be changed from the scenario editor.

b. Accessions

Accessions in OSAM can occur based on a pre-established plan, produced externally and imported into the database. These plans specify how many officers to access into each community. Alternatively, the user can allow OSAM to determine the number of accession via the unconstrained and constrained accessions options. The unconstrained accessions method calculates requirements for new accessions based on future requirements as projected from current inventory. Constrained accessions follow the same methodology, but prevent the addition of officers in excess of what OPA allows for the accessed paygrades and designators.

c. Promotions

OSAM can process promotions according to four different methods: promotion plan, promote-to-vacancy, promote-to-flowpoint, and auto-promote. The plan option is based on annually updated quotas of specific numbers of officers to promote in each community. The promote-to-vacancy option causes promotions to occur according to the number of losses that occur in the next higher paygrade, thus creating space for the lower-ranking officers to move into. The promote-to-flow-point option also tracks vacancies to manage promotions, but prevents officers from promoting earlier than what DOPMA allows. Along with the constraints in promote-to-flow-point, auto-promote restricts promotion opportunity to within DOPMA guidelines.

d. Transfers

Transfers between communities in OSAM occur according to prescribed plans developed outside the model. The scenario editor does not provide for changes to these plans and no other options are available. Transfer plans specify the number of in-quotas for receiving community designators and paygrades. Transfers out of communities follow a default system coded within OSAM that is based on historical distributions.

e. Losses

As discussed above, OSAM separately processes three types of losses: Natural, Force-Out, and User-Added losses. The user may adjust Natural losses by selecting different historical loss rate plans corresponding to varying levels of retention. Settings for Force-Out losses are not part of the scenario editor, but the user may change the criteria by modifying the code. The Force-Shaping tab on the scenario editor allows the user to enter in additional, specific numbers of losses for selected groups of officers.

G. RESULTS AND ANALYTICAL TOOLS

The run time for a single simulation in OSAM depends on the length of the user-specified time horizon and the speed of the computer. A typical five-year scenario takes approximately ten minutes on current desktop computers. Once the simulation completes, OSAM closes Access and the results are stored in an output database. The package of software included with OSAM contains a Scenario Analysis Tool developed in Microsoft Excel that reads in the results and can create charts, graphs, and reports for multiple scenarios. The Scenario Analysis Tool is effective for comparing individual simulations, but does not have the capacity to analyze a large-scale experiment with multiple replications.

III. EXPERIMENT DESIGN AND IMPLEMENTATION

This chapter discusses the use of Design of Experiments (DOE) to grow a dataset representing a broad range of scenarios. The process is analogous to farming: the seeds must be planted, the crops grow and are harvested, then the produce is cleaned, sorted, and turned into a finished product. In this study, data farming in OSAM begins with selecting the right factors and ranges, developing a design, and then making necessary arrangements to conduct the simulation. The SEED Center at the Naval Postgraduate School (https://harvest.nps.edu) provides the computing resources to run the experiment on a reasonable timeline.

A. DESIGN OF EXPERIMENTS

Applying DOE to simulations provides benefits that align with the purpose of OSAM and manpower planning in general. End-strength alone as a single number needs proper context to provide meaningful information to decision makers. DOE enables us to better understand the system in which those end-strength numbers arise and to explore the effect of potential policy changes on those systems (Kleijnen et al., 2005). By running a set of scenarios that encompass a wide range of possibilities, we gain a better understanding of (1) the model, (2) the way end-strength responds to changes in the factors varied in the experiment, and (3) the variability inherent in the responses. Unlike with other uses of DOE, the intent of this study is not to find an optimal set of settings to achieve a certain result, but rather to determine what conditions, as simulated by varying the settings, produce results that may warrant changes to policy. This study seeks to identify the most severe retention outcomes and examine the causes behind those particular results. Forces driving poor retention may often lie outside the control of the Navy; thus in order to ameliorate the disadvantages of taking a reactive approach, identifying scenarios requiring varying degrees of intervention takes on critical importance.

B. FACTOR SELECTION

In DOE terminology, the term “factor” refers to what may be called parameters, variables, or inputs for the model. The choice of factors for the design depends on the intent of the experiment and characteristics of the available factors in the model, and can also be constrained by the resources available to run the experiment (Kleijnen et al., 2005). Given the purpose of this study to exploit as fully as possible the mechanics of the model and explore retention, losses emerged as the primary category of inputs. Losses represent the most complex part of OSAM’s computations and they have the most significant impact on end-strength, which in turn provides metrics on retention. With three types of losses to choose from, the deciding considerations were ease of implementation and relationship to retention. User-Added losses provided the best choice to serve as factors in this experiment. The other two options, Natural Losses and Force-Outs, require manipulation of the tables within the OSAM database and changes to the code to allow the variations in a designed experiment. Changing historical loss rates that govern natural losses or changing criteria for Force-Outs also do not reflect a close relationship with retention. Historical loss rates account for a variety of influences that affect retention, so it would be difficult to determine what proportion is steady-state attrition and what proportion is caused by a changing retention environment. This confounding of effects would also hinder determining a valid range of values to build the design.

While it would be possible to incorporate additional inputs such as promotions and accessions into the design, given the limited group of officers comprising the subject of this study, these other inputs would likely not provide additional insight. At paygrades O3 and above and for simulation time-horizons on the order of five years, accessions have little impact. Promotion settings do affect the simulation in this case; however, rules and policies already constrain promotions to such a narrow range of effects that varying them would provide comparatively little benefit. As layed out in Table 1, the design consists of 27 factors, the number of forced losses per paygrade per designator per year.

Values for O5 and O6 are combined to reduce the size of the design to make it more manageable.

Table 1. List of factors used to build design.

Projection Year	Designator	Paygrade	Factor	Projection Year	Designator	Paygrade	Factor	Projection Year	Designator	Paygrade	Factor
1	1110	3	F1	2	1110	3	F10	3	1110	3	F19
		4	F2			4	F11			4	F20
		5/6	F3			5/6	F12			5/6	F21
	1120	3	F4		1120	3	F13		1120	3	F22
		4	F5			4	F14			4	F23
		5/6	F6			5/6	F15			5/6	F24
	1310/1320	3	F7		1310/1320	3	F16		1310/1320	3	F25
		4	F8			4	F17			4	F26
		5/6	F9			5/6	F18			5/6	F27

Factors are number of forced losses for each paygrade, designator, and fiscal year, for three years in the simulation.

The next step in the design development consists of determining an appropriate range of values for each factor. Guidance on what constitutes poor retention typically contains only qualitative descriptions, thus introducing a degree of subjectivity into the design. This study defines poor retention as 2 to 10 percent of additional attrition beyond historical losses. Using this definition, specific numbers of losses corresponding to those rates are calculated based on current inventories for each designator and paygrade combination. As summarized in Table 2, the 2 and 10 percent values represent the range of forced losses for each factor. These values are intended to represent mild to severe degrees of poor retention.

Table 2. Range of forced losses calculated as a percentage of inventory.

			Losses		
Designator	Grade	Count	10%	6%	2%
1110	3	2592	259	156	52
	4	1105	111	66	22
	5	879	88	53	18
	6	414	41	25	8
1120	3	1180	118	71	24
	4	542	54	33	11
	5	313	31	19	6
	6	225	23	14	5
1310	3	3379	338	203	68
	4	1559	156	94	31
	5	1107	111	66	22
	6	405	41	24	8
1320	3	1152	115	69	23
	4	721	72	43	14
	5	515	52	31	10
	6	213	21	13	4

C. DESIGN SELECTION

The same set of considerations that guide factor selection also help in selecting an appropriate design. A number of choices along a spectrum of complexity give the designer flexibility in the approach, although time and computing resource remain a constraint (Kleijnen at al., 2005). The large number of factors in this study and wide range of values mean gridded designs would not prove effective. Only a coarse grid may be possible in this case, which would not provide sufficient depth of coverage of the response surface. A fine grid would not be practicable due to the size of the factor space. Latin hypercubes emerge as the best candidate since they provide an excellent compromise between resolution and efficiency. A distinguishing feature of Latin hypercubes is their space-filling property, which scatters the design points throughout the design space to efficiently capture as much of the possible range of scenarios as possible. These designs are well suited to studies in which gaining a better understanding of the response surface is a primary goal (Sanchez & Wan, 2012).

Further efficiency and improved space filling properties can be gained by using a nearly orthogonal Latin hypercube (NOLH) design (Cioppa and Lucas, 2007).

After selecting NOLH as the design, the factors and range of values go into a design blueprint available for download from the SEED Center website (https://harvest.nps.edu/software.html). The blueprint calculates the factor values for each design point. With 27 factors, the NOLH design totaled 257 design points and still fills the design space quite well. Figure 4 provides a scatterplot matrix for just a portion of the factors and demonstrates the comprehensive coverage.

Figure 4. Scatterplot matrix for NOLH design using six out of 27 factors.

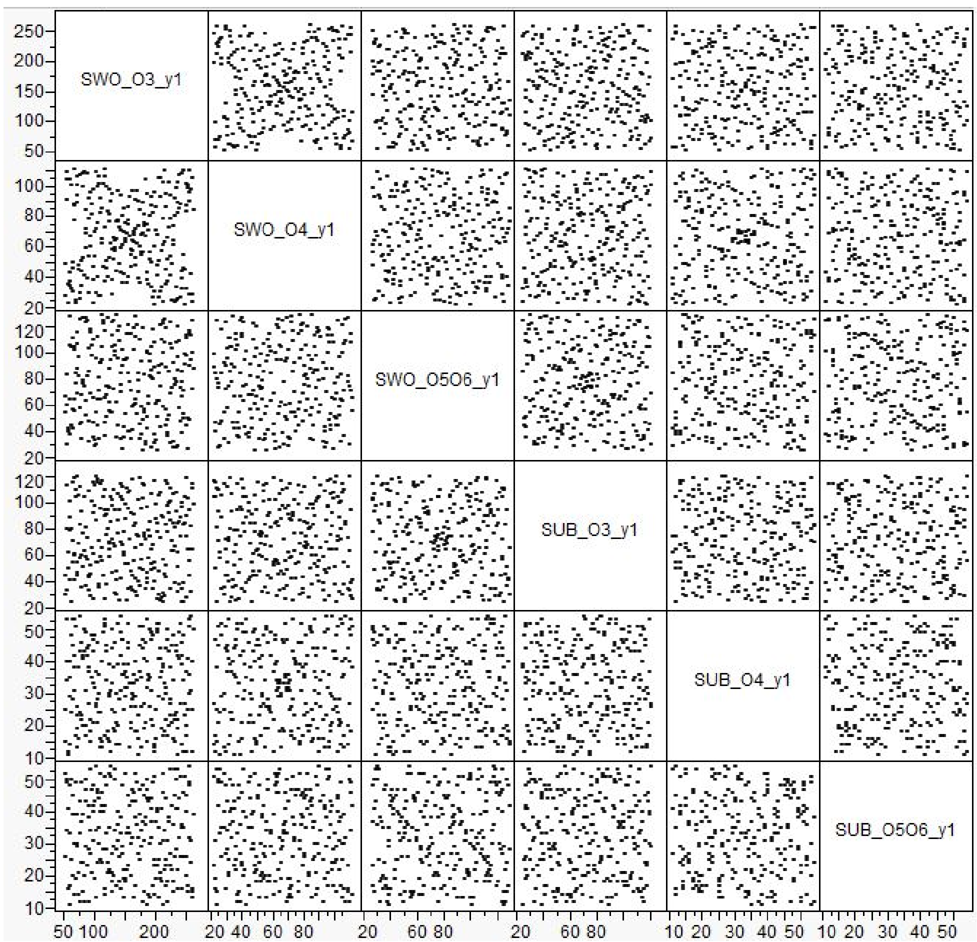

D. IMPLEMENTATION

Since OSAM does not have an organic capability to run multiple simulation runs successively, additional preparation was necessary to create the ability to run an experiment with multiple design points and replications. Software developed by Steve Upton from the SEED Center as part of this study can create designs based on user input, translate the design into OSAM input, and run the design in OSAM on a computer cluster.

1. OSAMFarmer

Once a user has selected a design and decided which factors to use, a user can run OSAMFarmer to draft the design. The program will vary the inputs to create the design points and will build all the necessary initialization files that OSAM uses to adjust the scenario settings for the simulation. In this study, since the design used an existing NOLH blueprint, running OSAMFarmer was not required. An additional script written in R translated the design from the NOLH spreadsheet into initialization files for OSAM.

2. OSAMRunner

The ability to run a series of simulations in OSAM without user intervention adds immense capability to the model. It allows a user to run large experiments, to replicate simulations with different random seeds to determine stochastic variation, or to automate the running of a list of scenarios not necessarily organized under a design framework. OSAMRunner, an executable program written in C++, replaces the scenario editor and allows running of OSAM from a command line interface. Once initialization files are placed in the correct location and random seeds are set, a multiscenario design can be run without interruption. Since each experiment still runs separately, multiple design points may be executed simultaneously, thus taking advantage of parallel processing capabilities with cluster computing.

3. Base Case

In addition to the 257 design points of the experimental design, this study also includes a single scenario to serve as a baseline for comparison experiment results. The base case contains the same parameters as the design points, but without the additional forced losses during the first three years of the simulation. These parameters include the 2014 accession and transfer plans, 2014 historical loss rates, and promotions by auto-promote. The auto-promote option sets promotion opportunity rates to follow predicted vacancies and DOPMA regulations on flow points representing years of commissioned service. To explore the stochastic variation more deeply, we replicate the base case scenario 100 times compared to ten replications for the remaining design points.

THIS PAGE INTENTIONALLY LEFT BLANK

IV. RESULTS AND ANALYSIS

This chapter contains a description of the results obtained from applying data farming to OSAM, analysis of those results, and metamodeling of the underlying systems. After assessing the degree of stochastic variation in the simulation, the averages of the replications are used as observations to build a set of linear regression models.

A. ANALYTICAL TOOLS

Organization of the data and analysis were completed using JMP Pro Version 10.0.0 and R version 3.2.0 (JMP Pro, 2013; R Core Team, 2015. Using filtering tools in JMP, we created data tables from the initial output. The data tables were loaded into R for various exploratory analyses and initial models. We used JMP to create the final models to take advantage of the graphical tools offered by the software.

B. OVERVIEW OF RESULTS

Using SEED Center cluster computers, the experiment took 30.6 hours to complete (S. Upton, personal communication, July 30, 2015). The initial output file from the simulation runs contains over 700 megabytes of data. These data included end-strength projections for all 68 designators for up to six paygrades over six fiscal years producing a total of 5,162,990 observations given the 257 design points and ten replications per design. In JMP, the data are pared down to only the four designators of interest in this study, 111X Surface Warfare, 1120X Submarine Warfare, 1310 Aviation Warfare-Pilot, and 1320 Aviation Warfare-Naval Flight Officer (NFO).

C. STOCHASTIC VARIATION

OSAM produces only a small amount of stochastic variation in the results. Across all design points and the base case, the results contained an almost negligible amount of variation across replications. After 100 replications of the base-case scenario, the results for total end-strength of the URL communities examined in this study remain tightly clustered. Summary statistics for these replications are provided in Table 3. Stochastic variation across the ten replications of the design points similarly remained very low. From the plots of end-strengths for a representative sample of design points and officers groups, given in Figures 5 through 8, the results stay consistent across the replications represented by individual lines in the plots. Variance appears to increase slightly in later years in the simulation.

Table 3. Summary statistics of 100 replications of base case scenario, by fiscal year.

	2015	2016	2017	2018	2019	2020
Minimum	22880	23080	23270	23510	23670	23770
1st Quartile	22890	23110	23320	23540	23720	23830
Median	22890	23120	23330	233550	23740	23850
Mean	22890	23120	23330	23550	23740	23840
3rd Quartile	22890	23130	23350	23570	23750	23860
Maximum	22900	23160	23390	23580	23790	23900

Figure 5. Line graph of end-strength for O3 Submarine Officers for design point 1.

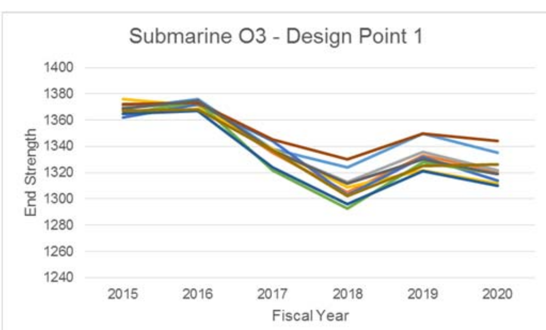

Figure 6. Line graph of end-strength for O4 SWOs for design point 20.

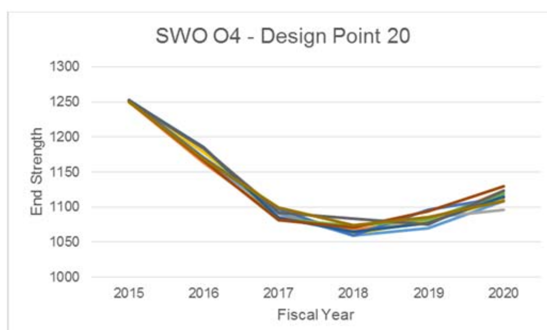

Figure 7. Line graph of end-strength for O5 NFOs for design point 100.

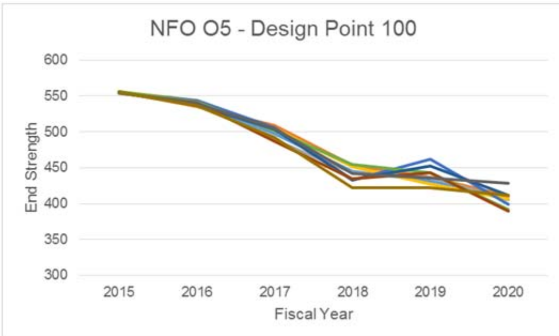

Figure 8. Line graph of end-strength for O6 Pilots for design point 200.

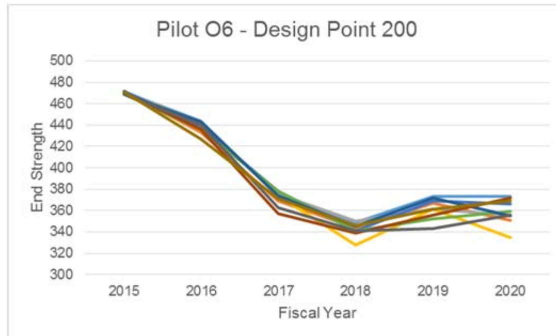

D. DESCRIPTIVE STATISTICS

OSAM produces several pieces of data such as numbers of losses and transfers during the course of a simulation run, but since it is an inventory projection model, the end-strength represents the end-product of the separate calculations.

The multidimensional aspect of the inputs creates a dataset with a corresponding degree of complexity. Although some insights may be gained by analyzing the entire officer corps as a whole, this study focuses on specific communities, identified by designator and paygrade, to determine the effects of losses on retention. Since we do not know beforehand how the response may vary across groups, we first examine each group’s results separately. Appendix A contains box plots of the end-strengths for all sixteen groups. See Figure 9 below for explanation of how the descriptive statistics are arranged in a box plot.

Figure 9. Explanation of box plot structure (T. Lucas, personal communication, 2014).

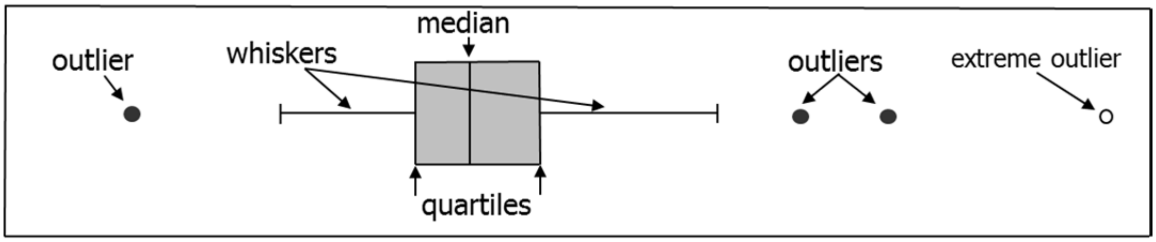

From these plots, we identify some patterns that remain consistent across all paygrades. Box plots of the summarized end-strengths by paygrade in Figures 10 through 13 by paygrade show how the variance changes by fiscal year, and the trends in end-strength changes. In each plot, the results for the first year of results shows very little variance. A possible explanation could be that the forced losses introduced in the design have not yet had a chance to exert a strong influence on end-strength. The promotion policies in place to manage inventory may create a system that is robust enough to cope with the smaller number of forced losses introduced for just one year. By the following year, a clear effect is evident both in the lower end-strength and variation of the results.

Figure 10. Box plot of O3 end-strength.

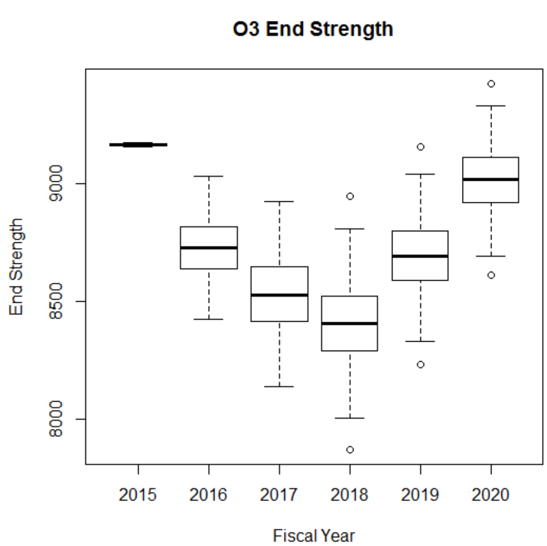

Figure 11. Box plot of O4 end-strength.

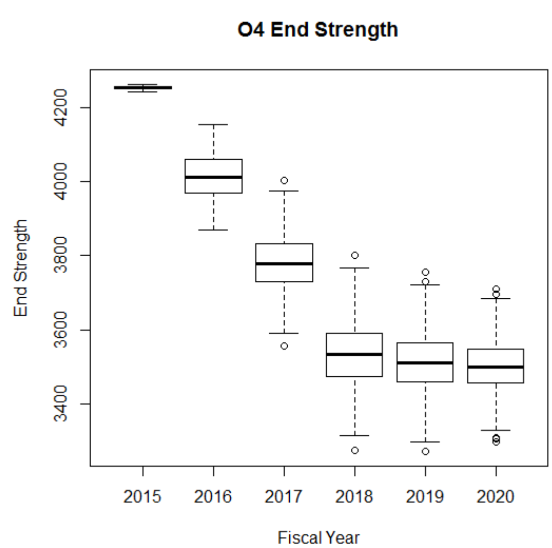

Figure 12. Box plot of O5 end-strength.

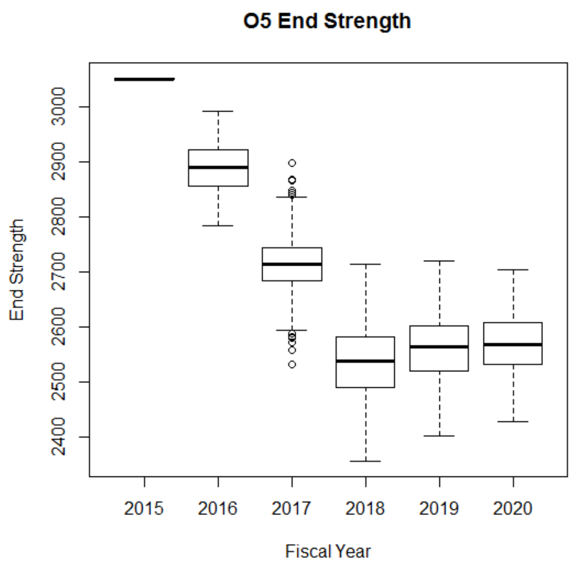

Figure 13. Box plot of O6 end-strength.

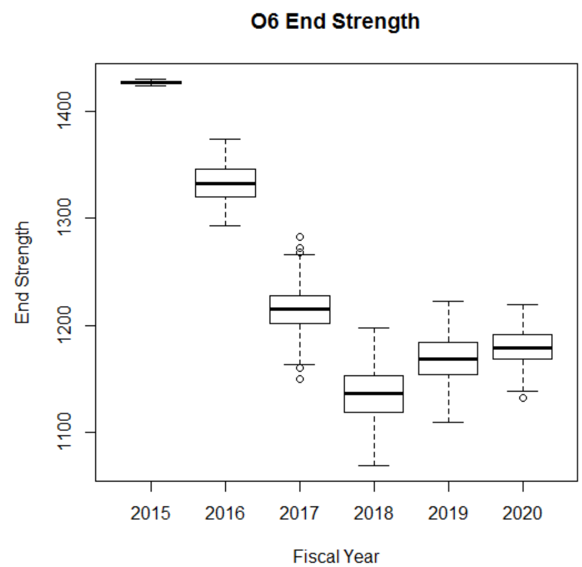

For junior and mid-grade officers (O3 and O4), the range variance appears to increase through 2018 and then remain stable through the remaining two years. For senior officers (O5 and O6), the variation peaks in 2018 and then decreases slightly. The end-strength for O3s displays the most resilience. Although the forced losses cause a significant decrease of over 700 officers, the end-strength quickly recovers once the additional losses end. Within two years after reaching minimum, average O3 end-strength is close to the 2015 value. O4, O5, and O6 end-strengths appear much less resistant to the simulation period of poor retention. End-strength in these paygrades increases only slightly, if at all, through the remaining years in the simulation.

E. RESPONSE VARIABLES

As the response variable of interest in this study, the end-strength must have a frame of reference to give meaning to any metamodels. To gauge the effect of poor retention on the health of the force, we compare the results of the designed experiment with the base case. End-strengths from the experiment represent the whole landscape of poor retention scenarios while the base case represents an unaffected scenario. In analyzing the results, we use the deviation of end-strength from the base case value as the response variable. Within the data table containing the results, these values are calculated by subtracting the design point end-strength from the corresponding base case end-strength for each fiscal year, paygrade, and designator combination. Design point and base case values for these calculates are the averages across the replications. Positive values in the deviation indicate that the design point end-strength is less than the respective base case value. Table 4 provides a summary of the means and standard deviations for each group for the three fiscal years that make up the post-poor retention period.

Table 4. Descriptive statistics of deviation from end-strength.

		2018		2019		2020	
Grade	Designator	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
O3	SWO	943	74	1041	58	1095	46
	SUB	430	33	465	25	475	16
	PILOT	948	103	901	87	824	76
	NFO	332	45	337	32	325	25
O4	SWO	104	41	92	45	94	53
	SUB	62	21	85	23	69	26
	PILOT	181	56	161	56	153	51
	NFO	92	21	58	19	25	18
O5	SWO	146	33	122	30	97	31
	SUB	37	13	11	13	26	13
	PILOT	52	40	65	35	48	35
	NFO	41	19	42	17	58	15
O6	SWO	30	18	18	20	-2	19
	SUB	4	10	0	9	-8	9
	PILOT	82	17	66	18	83	15
	NFO	12	9	10	8	9	8

Means and standard deviations of the response variable used to build metamodels: the deviance of the design point end-strength from the base case end-strength.

Over the three years, means both increase and decrease by varying degrees for different officer groups. No clear trends or patterns emerge for how much the simulation varied from the base case when broken down by paygrade and designator. With a few exceptions, however, the standard deviations show a consistent decreasing trend.

1. Distribution of the Response Variable

Examining the distribution of the response may reveal further insights into potential central tendency. We classify the distribution based on the histograms provided in Figures 14 and 15. These histograms represent end-strengths summed across the four paygrades and four designators for their respective fiscal years. Histograms of end-strengths for individual officer groups displayed similar shapes. The distributions of outcomes appears to be roughly normal. Distinct peaks in both histograms represent some deviation from a smooth normal curve.

Figure 14. Histogram of 2018 end-strengths across all paygrades and designators.

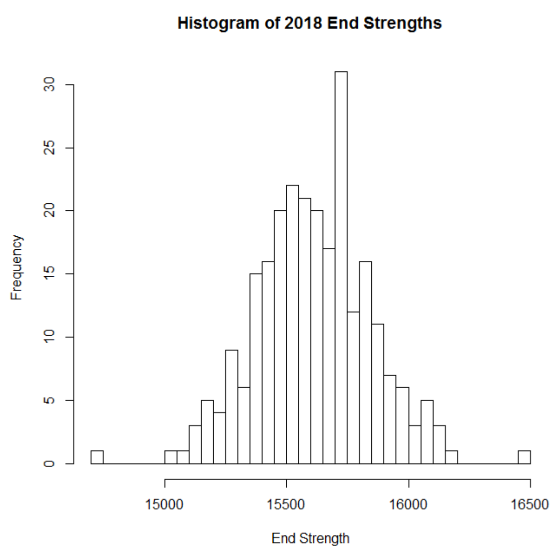

Figure 15. Histogram of 2020 end-strengths across all paygrades and designators.

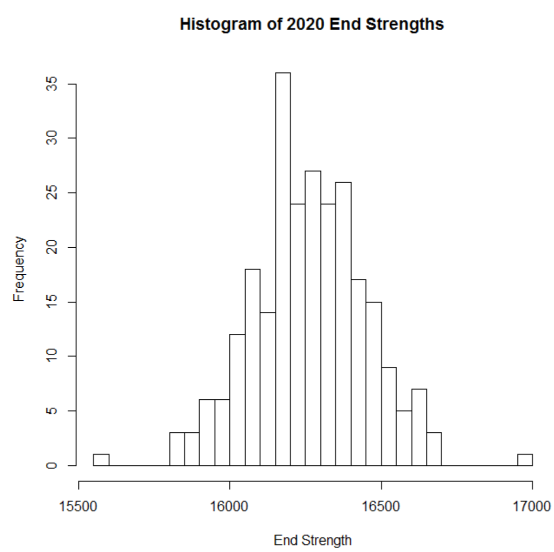

2. Correlation between Officer Groups

Given that the response variable, deviation of the design point end-strength from the base-case end-strength, can be broken down by paygrade, designator, and fiscal year, we may consider using multivariate multiple regression instead of having multiple linear regression models. Multivariate multiple regression provides advantages when the responses are highly correlated. The pairwise plots of the Pilot and NFO end-strengths in Figure 16 show a high degree of correlation. The likely reason for this result is the treatment of Pilots and NFOs as one group in the experimental design. Forced losses were applied evenly to all aviators. As evident in the plots in Figure 17, SWO and Submarine officer results did not exhibit correlation between each other, nor when compared to aviation officers.

Figure 16. Scatterplot of Pilot and NFO end-strengths.

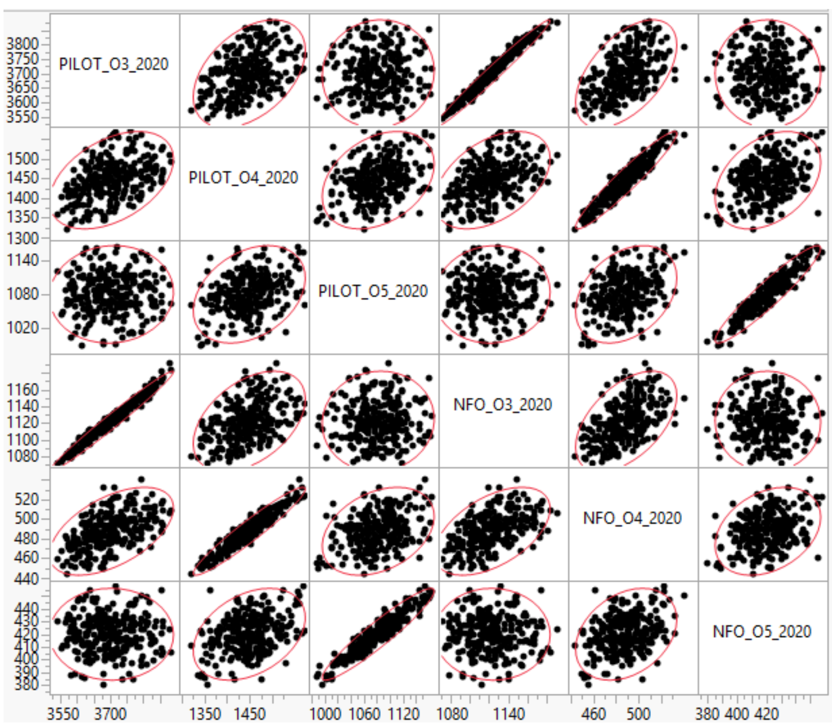

Figure 17. Scatterplot of SWO, Submarine Officer, and Pilot end-strengths.

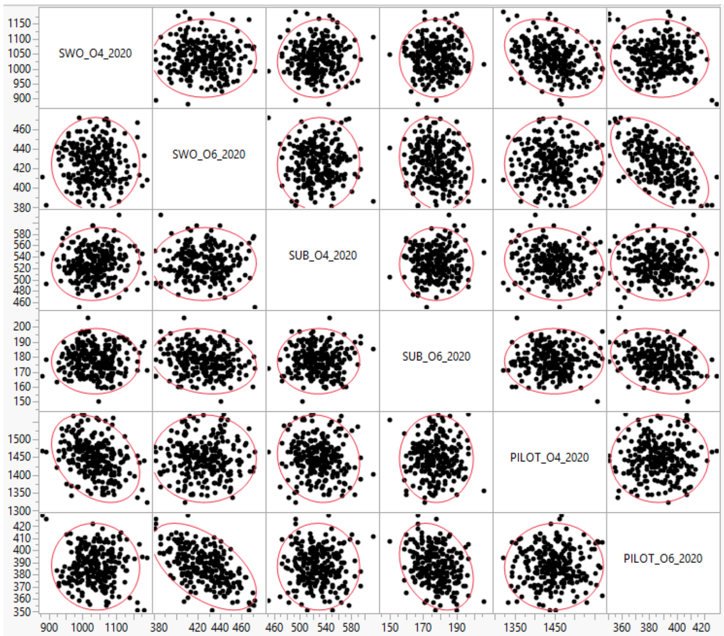

F. METAMODELS

To the manpower planner attempting to forecast end-strength based on loss data, having a valid model of the relationship between loss and end-strength is a critical component of the analysis. Multiple options exist for model types and even more choices come with the decisions on how to build the model. Consequently, the scope of possibilities for how to model the relationship between end-strength and losses presents another opportunity for structured quantitative analysis to assist in the process. In this section, we present a few candidate models that attempt to maximize accuracy and usefulness to the planner. Observations are organized into the 27 predictor variables listed in Table 5. Each variable represents the amount of forced losses for different groups in each year of the three-year period of poor retention. O5 and O6 losses were combined in the design, as were losses Pilots and NFOs.

Table 5. List of 27 predictor variables in dataset.

Predictor Variables		
SWO_O3_y1	SUB_O3_y1	AV_O3_y1
SWO_O3_y2	SUB_O3_y2	AV_O3_y2
SWO_O3_y3	SUB_O3_y3	AV_O3_y3
SWO_O4_y1	SUB_O4_y1	AV_O4_y1
SWO_O4_y2	SUB_O4_y2	AV_O4_y2
SWO_O4_y3	SUB_O4_y3	AV_O4_y3
SWO_O5O6_y1	SUB_O5O6_y1	AV_O5O6_y1
SWO_O5O6_y2	SUB_O5O6_y2	AV_O5O6_y2
SWO_O5O6_y3	SUB_O5O6_y3	AV_O5O6_y3

1. Multivariate Analysis

Experiments with more than one response variable may be analyzed using multivariate methods that build multiple models simultaneously. Although such methods may apply to this study, we build separate linear models for each response to reflect the paygrade- and designator-specific approach to manpower planning used by the Navy.

2. Multivariate Regression Models Based on One Year Loss Data

A model that could detect leading indicators of a poor retention trend would have significant value to planners. Reasonably accurate projections of end-strength based on one year of loss data and assumptions about the retention trend, would allow planners to take a proactive approach to the problem. Taking the numbers of forced losses for the first year of the simulation, we attempt to build linear models with end-strength of the fourth year, 2018, as the response. These models require an assumption that the poor retention period will last three years. The linear models were built in using the lm() function in R with 16 models total, one for each paygrade and designator combination. To keep the models as simple as possible, stepwise regression was performed using the stepAIC() function from the MASS package (Venables and Ripley, 2002). This function removes unnecessary terms from the model based on the Aikaike Information Criterion (AIC) to assess candidate models. The R-squared values, summarized in Table 6, show that these models perform very poorly; therefore further analysis was conducted. For a full explanation of linear models, Faraway's text (2005) provides a useful reference.

Table 6. R-squared of end-strength projection models based on one year of loss data.

Grade	Designator	R^2
O3	SWO	0.162
	SUB	0.119
	PILOT	0.240
	NFO	0.319
O4	SWO	0.301
	SUB	0.421
	PILOT	0.250
	NFO	0.276
O5	SWO	0.295
	SUB	0.273
	PILOT	0.175
	NFO	0.245
O6	SWO	0.187
	SUB	0.135
	PILOT	0.227
	NFO	0.197

3. Multivariate Regression Models Based on Three Years of Loss Data

Using all three years of the forced losses makes the maximum number of predictor variables available to the model. For the response variable, we use the 2020 end-strength. The validity of this structure depends on the assumption that the poor retention period lasts three years followed by three years of losses in line with historical rates. Although a user would not be able to predict with certainty when the poor retention period may end, planners could use models of this type to conduct what-if analyses with a three-year time horizon.

Since officer inventory is managed by community and policies are often applied differently by paygrade, we create one model for each paygrade and designator combination. Out of these 16 models, we present the results of a representative sample of four models representing each paygrade and designator. More comprehensive results are provided in Appendix B. The four groups presented here are Submarine Warfare O3s, SWO O4s, Pilot O5s, and NFO O6s.

A common approach to building linear models is to start with a wide scope and include all predictor variables as well as possible interaction and nonlinear terms (Crawley, 2013). In this study, initial models include all main effects consisting of the 27 loss data variables and all two-way interactions. Following the principle of parsimony, we use stepwise regression to eliminate unnecessary terms. Using the Stepwise Regression Control option in JMP, we select P-value Threshold as the stopping rule with a p-value of 0.01 for both the probability-to-enter and probability-to-leave criteria.

To ensure that the models are not overfit we assess how well each model performs with new data. We randomly partition the data into training and test sets that represent 80% and 20% of the data, respectively. Using a model built with only the training set, we compare the predicted values based on the data in the test set with the actual values for those observations.

a. Submarine Warfare Officers - O3

The initial model produced in JMP using all the data and the standardized approach contained 14 predictor variables and had an R^2 of 0.969. This model contained all nine loss variables for O3s plus two aviation loss variables for O4 and two SWO and aviation interaction terms. The aviation and interaction terms had coefficients of two to three orders of magnitude smaller values than the O3 terms, and had higher p-values. After accounting for the difference in units between the coefficients, we conclude that the interaction terms do not add value to the model and exclude them from the final model summarized in Table 7.

Table 7. Parameter estimates for 2020 end-strength of O3 Submarine Officers.

Term	Estimate	Std Error	t Ratio	Prob (> |t|)
Intercept	1406.379	1.544	910.7	<0.001
SWO_O3_y1	-0.044	0.003	-13.63	<0.001
SWO_O3_y2	-0.036	0.003	-11.16	<0.001
SWO_O3_y3	-0.036	0.003	-11.17	<0.001
SUB_O3_y1	-0.38	0.007	-5.34	<0.001
SUB_O3_y2	-0.233	0.007	-32.98	<0.001
SUB_O3_y3	-0.433	0.007	-61.19	<0.001
AV_O3_y1	-0.041	0.002	-22.05	<0.001
AV_O3_y2	-0.04	0.002	-21.21	<0.001
AV_O3_y3	-0.04	0.002	-21.03	<0.001

The R^2 of 0.964 shows that the model is highly accurate. Based on the size of the coefficients, the O3 Submarine Officer losses in years two and three have the largest effect on 2020 end-strength. A model built with only those two predictors had a R^2 of 0.698, indicating that the SWO and aviation losses still have a significant effect.

Diagnostic plots of the models indicate that key modelling assumptions are met. The residuals versus predicted plot in Figure 18 has a homoscedastic pattern and the normal Quantile-Quantile (Q-Q) plot in Figure 19 shows that the residuals have an approximately normal distribution.

Figure 18. Residuals versus predicted values of linear model for O3 Submarine Officers.

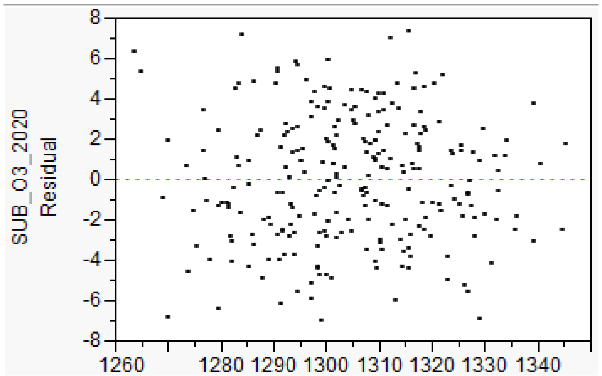

Figure 19. Normal Q-Q plot of the residuals for O3 Submarine Officer model.

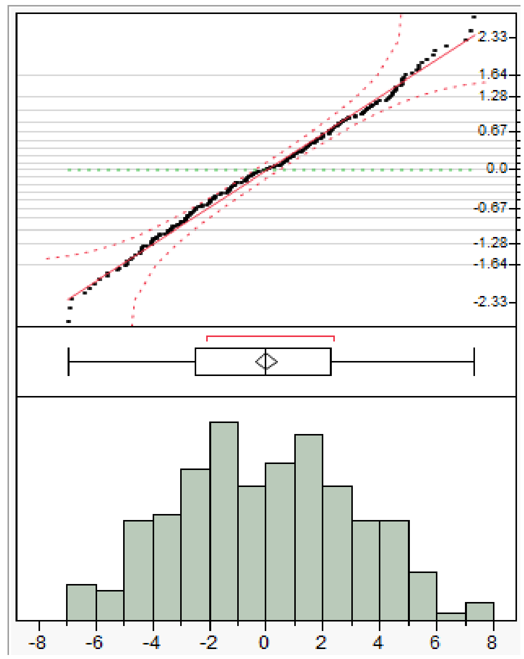

Although the final model eliminates unnecessary terms, the high R^2 indicates that the model could be overfit. We rebuild the model using only the training set and compare the predicted results with the actual values from the test set observations. The plot of actual versus predicted values, provided in Figure 20, shows a diagonal linear shape indicative of fairly accurate performance.

Figure 20. Actual versus predicted plot for test data set using O3 Submarine Officer model.

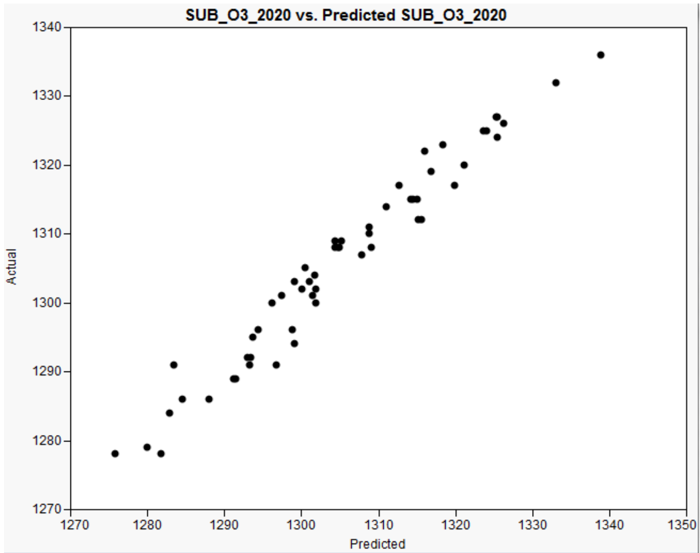

b. Surface Warfare Officers – O4

For O4 SWOs, a stepwise regression model of all main effects and two way interactions produced a model with 26 predictor variables with 18 main effect variables for O3s and O4s and eight interaction terms, four of which involved only aviation groups. This model has an R^2 of 0.984. We remove the interactions terms since they contribute very little to the model based on the small coefficients. Our final model still maintains a high R^2 of 0.971. SWO O3 and SWO O4 variables have the greatest influence. A model built with only O4 losses for all designators has an R^2 of 0.247 and a model built with only SWO losses for O3 and O4 has an R^2 of 0.545. Parameter coefficients and their p-values of the final model are provided in Table 8.

Table 8. Parameter estimates for 2020 end-strength of O4 SWOs.

Term	Estimate	Std Error	t Ratio	Prob (> |t|)
Intercept	1209.257	6.532	185.13	< 0.0001
SWO_O3_y1	-0.347	0.01	-35.33	< 0.0001
SWO_O3_y2	-0.353	0.01	-35.99	< 0.0001
SWO_O3_y3	-0.284	0.01	-28.97	< 0.0001
SWO_O4_y1	-0.331	0.023	-14.46	< 0.0001
SWO_O4_y2	-0.403	0.023	-17.59	< 0.0001
SWO_O4_y3	-0.52	0.023	-22.68	< 0.0001
SUB_O3_y1	0.149	0.021	6.99	< 0.0001
SUB_O3_y2	0.152	0.021	7.16	< 0.0001
SUB_O3_y3	0.128	0.021	6.06	< 0.0001
SUB_O4_y1	-0.159	0.048	-3.35	0.0001
SUB_O4_y2	-0.138	0.048	-2.91	0.004
SUB_O4_y3	-0.156	0.048	-3.29	0.001
AV_O3_y1	0.146	0.006	25.97	< 0.0001
AV_O3_y2	0.175	0.006	31.03	< 0.0001
AV_O3_y3	0.152	0.006	26.95	< 0.0001
AV_O4_y1	-0.19	0.011	-17.09	< 0.0001
AV_O4_y2	-0.224	0.011	-20.14	< 0.0001
AV_O4_y3	-0.173	0.011	-15.53	< 0.0001

In verifying the model assumptions, the residual versus predicted plot in Figure 21 of the residuals versus predicted values shows a homoscedastic pattern and the normal Q-Q plot in Figure 22 indicates that the residuals have a somewhat skewed, but still roughly normal, distribution. A plot of predicted versus actual values for the test set in Figure 23 indicates good performance for new data.

Figure 21. Residuals versus predicted values of linear model for O4 SWOs.

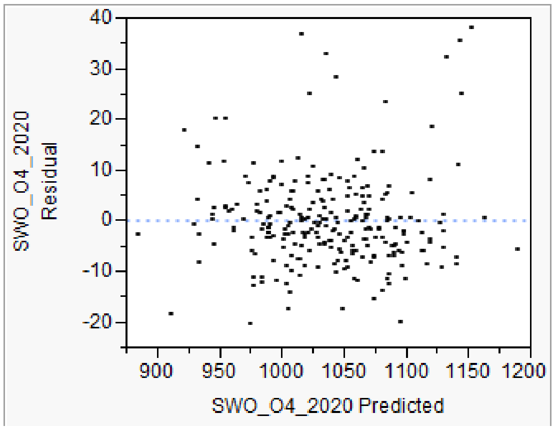

Figure 22. Normal Q-Q plot of the residuals for O4 SWOs.

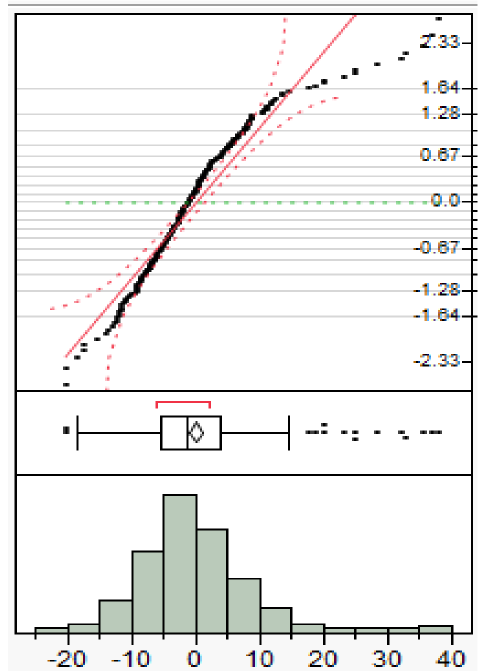

Figure 23. Actual versus predicted plot for test set using O4 SWO model.

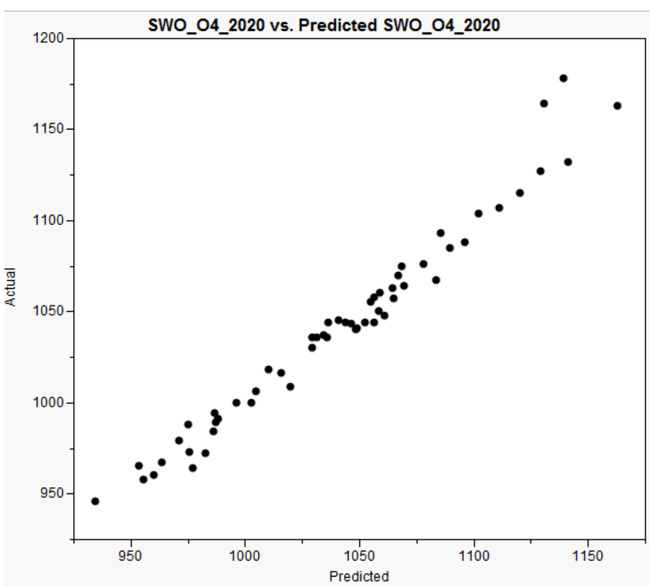

c. Pilots – O5

The initial stepwise regression model for O5 pilots contains 23 terms, including 20 main effects variables and three interactions terms, and has an R^2 of 0.980. The inclusion of interaction terms in the regression process yields very little benefit. Rerunning the stepwise regression with only main effects produces a model with 18 variables and an R^2 of 0.974. Table 9 contains a summary of parameters for the final model.

Table 9. Parameter estimates for 2020 end-strength of O5 Pilots.

Term	Estimate	Std Error	t Ratio	Prob (> |t|)
Intercept	1283.657	4.039	317.78	< 0.0001
SWO_O3_y1	0.016	0.006	2.66	< 0.0001
SWO_O4_y1	0.116	0.014	8.15	< 0.0001
SWO_O4_y2	0.152	0.014	10.75	< 0.0001
SWO_O4_y3	0.159	0.014	11.22	< 0.0001
SWO_O5O6_y1	-0.086	0.012	-7.02	< 0.0001
SWO_O5O6_y2	-0.115	0.012	-9.39	< 0.0001
SWO_O5O6_y3	-0.103	0.012	-8.4	< 0.0001
SUB_O4_y2	0.156	0.029	5.3	< 0.0001
SUB_O4_y3	0.116	0.029	3.96	< 0.0001
SUB_O5O6_y1	-0.083	0.029	-2.81	0.005
SUB_O5O6_y2	-0.093	0.029	-3.17	0.002
AV_O3_y1	-0.034	0.003	-9.76	< 0.0001
AV_O4_y1	-0.162	0.007	-23.47	< 0.0001
AV_O4_y2	-0.231	0.007	-33.55	< 0.0001
AV_O4_y3	-0.211	0.007	-30.57	< 0.0001
AV_O5O6_y1	-0.273	0.007	-39.17	< 0.0001
AV_O5O6_y2	-0.3	0.007	-43.08	< 0.0001
AV_O5O6_y3	-0.33	0.007	-47.37	< 0.0001

Diagnostic plots confirm that basic modelling assumptions are met. The plot of residuals versus predicted values in Figure 24 shows homoscedasticity and the normal Q-Q plot in Figure 25 shows that the residuals are normally distributed. The model performs well in predicting results for the test set data as demonstrated by the diagonal linear shape in the plot of actual versus predicted values in Figure 26.

Figure 24. Residuals versus predicted values of linear model for O5 Pilots.

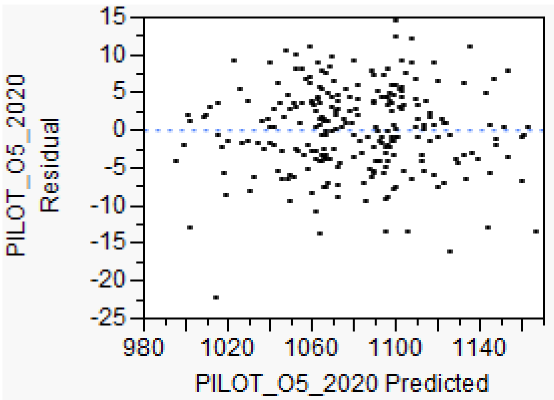

Figure 25. Normal Q-Q plot of the residuals for O5 Pilots.

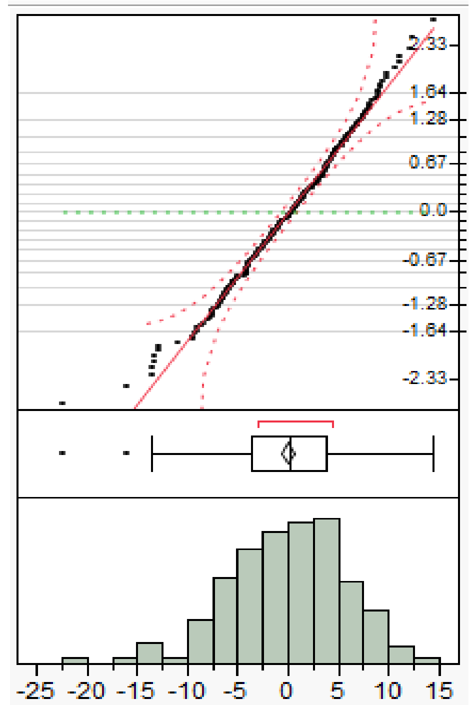

Figure 26. Actual versus predicted plot for test set using O5 Pilot model.

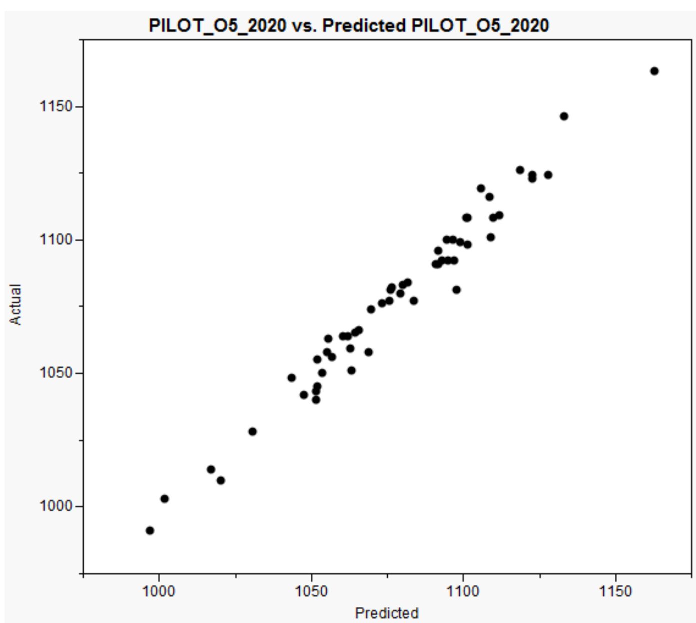

d. Naval Flight Officers – O6

Stepwise regression for the NFO O6 data did not include any interaction terms under the criteria used. The model used only the O5-O6 loss data with the addition of the aviation O4 losses for the first year. The final model had eight variables and an R^2 of 0.900, summarized in Table 10. The lower accuracy of the model may partly be due to the smaller population size of the officer group compared to groups in other paygrades and designators.

Table 10. Parameter estimates for 2020 end-strength of O6 NFOs

Term	Estimate	Std Error	t Ratio	Prob (> |t|)
Intercept	211.648	1.134	86.59	< 0.0001
SWO_O5O6_y1	0.033	0.005	6.31	< 0.0001
SWO_O5O6_y2	0.04	0.005	7.67	< 0.0001
SWO_O5O6_y3	0.034	0.005	6.6	< 0.0001
SUB_O5O6_y1	0.037	0.012	3.02	0.003
AV_O4_y1	-0.008	0.003	-2.81	0.005
AV_O5O6_y1	-0.061	0.003	-20.79	< 0.0001
AV_O5O6_y2	-0.079	0.003	-27.01	< 0.0001
AV_O5O6_y3	-0.089	0.003	-30.18	< 0.0001

The diagnostic plots in Figures 27 and 28 verify agreement with modelling assumption as seen in the other models. The plot of actual versus predicted values for test set data in Figure 29 exhibits linear behavior, although the performance does not appear as strong as compared to previous models. This result is consistent with the lower R^2.

Figure 27. Residuals versus predicted values of linear model for O6 NFOs.

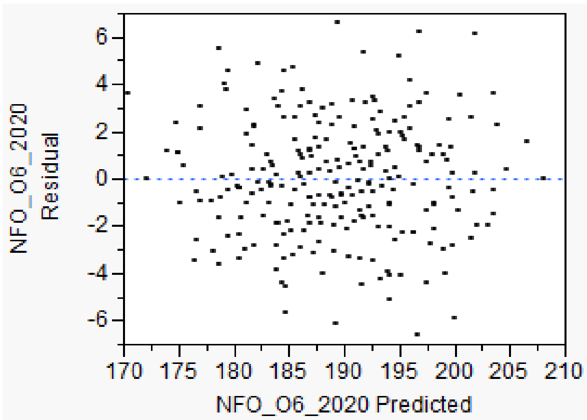

Figure 28. Normal Q-Q plot of the residuals for O6 NFOs.

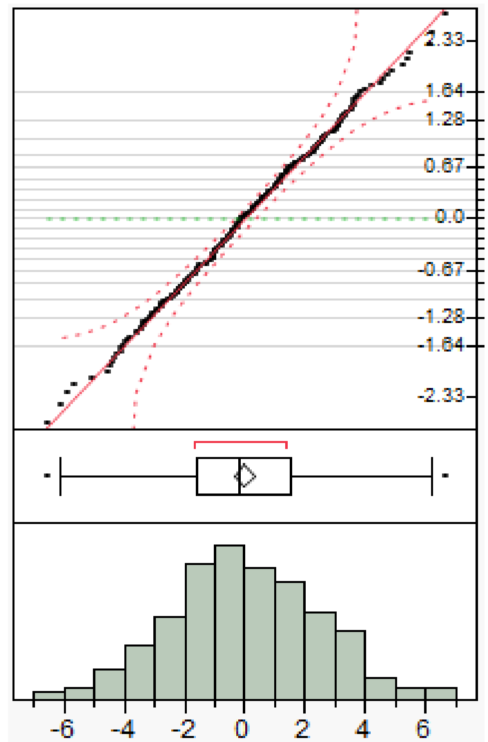

Figure 29. Actual versus predicted plot for test set using O6 NFO model

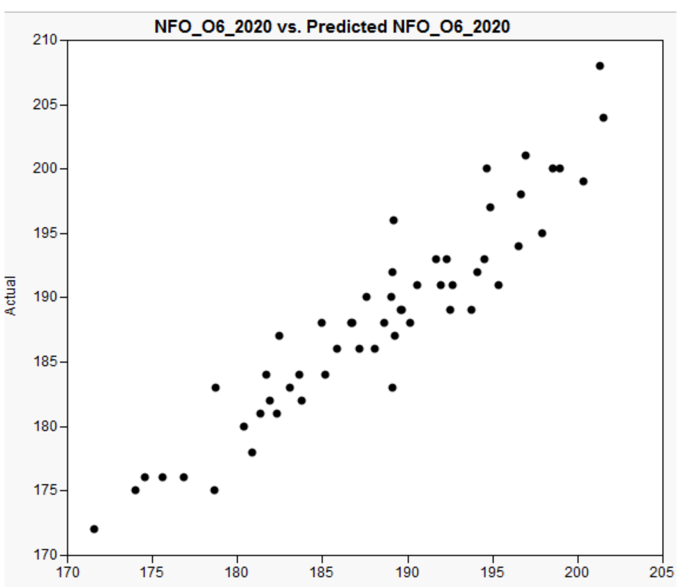

THIS PAGE INTENTIONALLY LEFT BLANK

V. CONCLUSIONS

A. ANSWERS TO RESEARCH QUESTIONS

1. Under What Conditions Do Officer Communities Experience Retention Problems?

The intent of this study is to provide a quantitative answer to this research question. As a reflection of how resilient officer communities may be to poor retention, the key issue is how many additional losses a particular group can sustain without severely degrading the long-term health of the inventory. Although temporary drops in inventory caused by a single year of high losses are acceptable provided that the community can recover, an end-strength reduction that persists more than two years can negatively impact readiness.

The results shown in the boxplots in Chapter IV and Appendix A show that some communities weather the poor retention better than others in the model. The complexity of the models indicates that outcomes depend not just on the losses occurring within the specific paygrades and designators, but also on the losses in other communities as well. For O4 SWOs, the end-strength after the six-year period of the simulation cannot be predicted based only on the additional losses experienced by SWOs only or O4s only. An accurate prediction depends on losses experienced by other communities as well. In this respect, models built using a complete representation of response surface can assist. The plots of SWO O4 end-strength in 2020 versus the forced losses in each year, provided in Figures 30 through 32 show that no single factor can reliably predict the outcome. In building the linear models, R^2 for individual groups decreased when removing variables for other groups, suggesting that the whole URL community must be considered when analyzing losses. The variation of models and results across the different officer groups make it difficult to identify precise conditions for a retention problem. The linear models built in Chapter IV provide an informative starting point when analyzing officer groups individually.

Figure 30. Plot of SWO O4 end-strength versus year one losses.

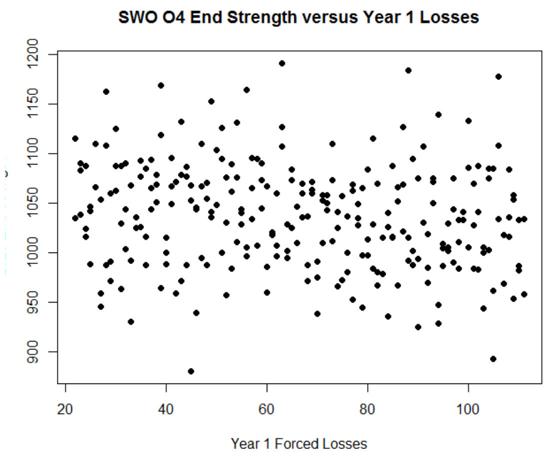

Figure 31. Plot of SWO O4 end-strength versus year two losses.

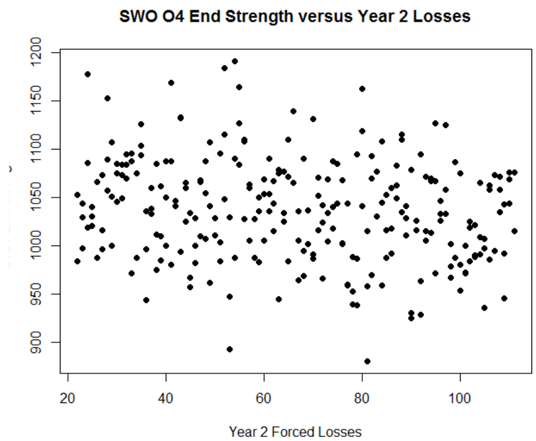

Figure 32. Plot of SWO O4 end-strength versus year three losses.

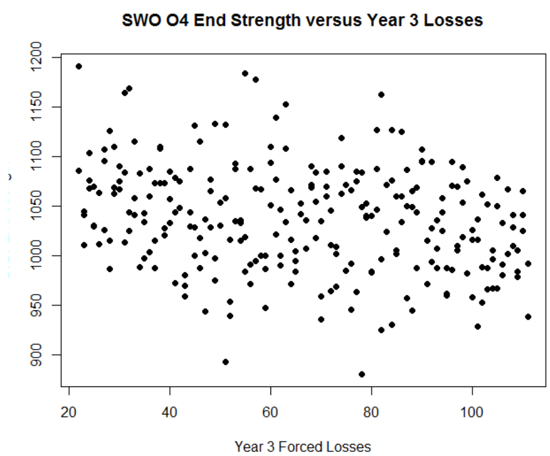

2. What Is the Impact of a Sustained Period of Poor Retention?

Individual groups reacted differently in the experiment. Overall, URL O3s recovered well from the losses within three years of the end of the poor retention period. Even under a worst-case scenario of high losses in each year, with a 7 percent total loss by year six, the end-strengths still follow a upward trend that suggest full recovery within one or two years. The losses sustained for O4s made a strong impact, which made it difficult for the group to recover after the period of losses ended. After a decline of approximately 14 percent, O4 end-strength remained steady at the reduced level for the remainder of the simulation. O5 and O6 officers suffered similar declines and fail to rebound by end of 2020.

3. How Much Does the Response Vary?

Variation in OSAM covers two areas: stochastic variation for a particular scenario and the variation of outcomes based on an experimental design. Stochastically, OSAM produces little variation when multiple replications are run with random seeds. The consistency of output justifies the current practice of using one replication per scenario.

Using the space-filing NOLH as the experimental design created a wide range of scenarios to run in OSAM. The corresponding outputs vary significantly as well, but still within reasonable ranges given the context. Differences in results between design points tend to fall within a range equal to approximately 10 percent of the end-strength.

B. RECOMMENDATION FOR FUTURE STUDY

The primary benefit of this study is that it provides the tools to conduct further data farming experiments in OSAM. Varying forced losses across three years for a particular set of officer groups represents just one of many possible applications. With the OSAMFarmer and OSAMRunner programs, users can expand on the scenario presented in this study and build new designs to explore other aspects of end-strength management beyond retention. To continue the experiments used here, we recommend running additional simulations that expand the parameters of the period of poor retention. Instead of having a standardized three-year period, further experiments could lengthen and shorten the period, and incorporate buildup and slowdown phases. Additional designators and groups of officers should also be studied to compare the impact of additional losses between communities. Varying other model inputs, such as promotions, transfer, and historical loss rates, may also yield further insight.

APPENDIX A. BOX PLOTS OF ALL DESIGNATORS AND PAYGRADES

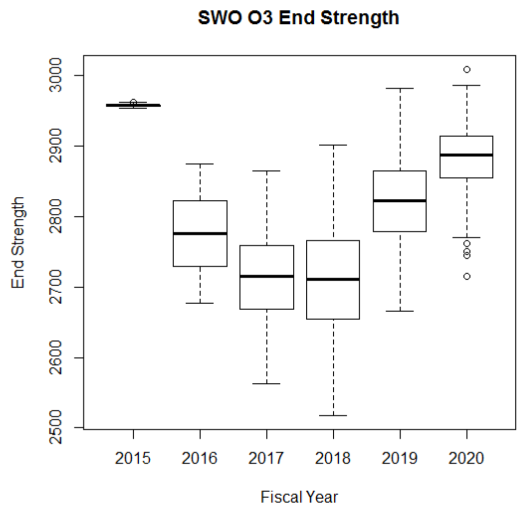

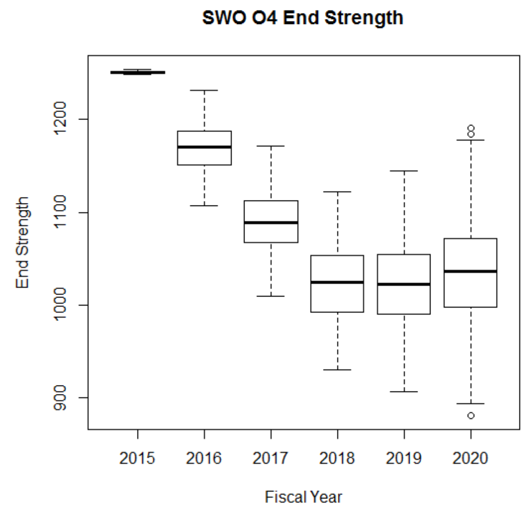

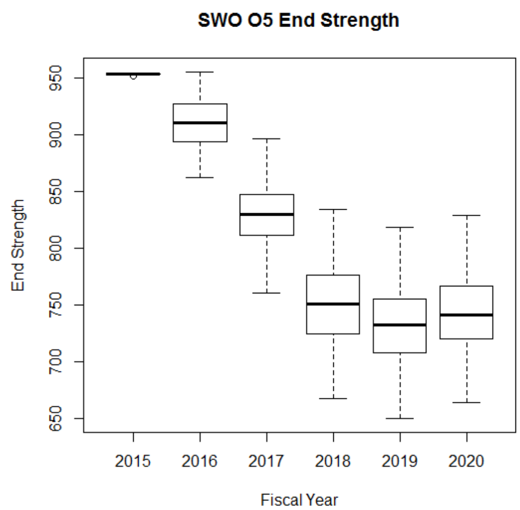

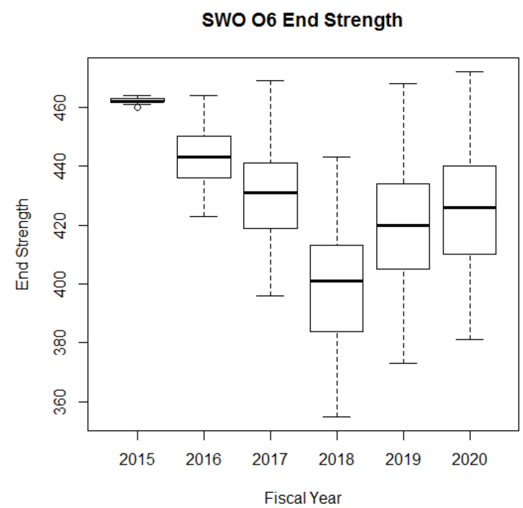

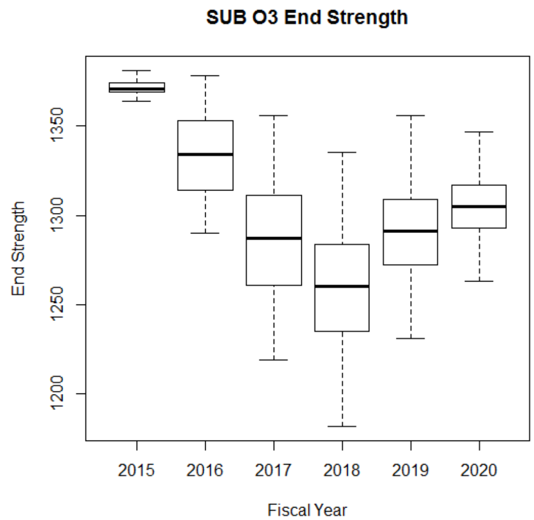
SUB O3 End Strength
End Strength
1200
1250
1300
1350
2015
2016
2017
2018
2019
2020
Fiscal Year

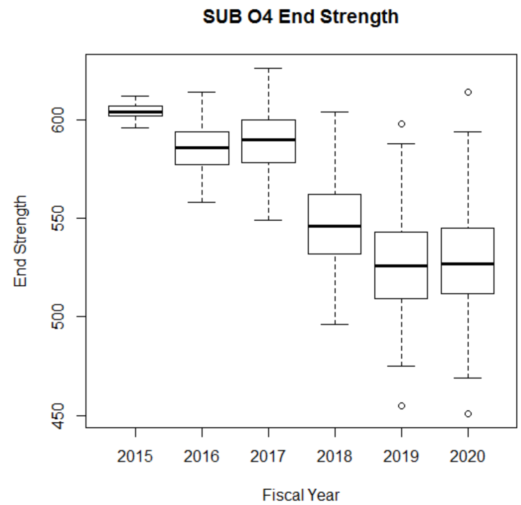
SUB O4 End Strength
End Strength
450
500
550
600
2015
2016
2017
2018
2019
2020
Fiscal Year

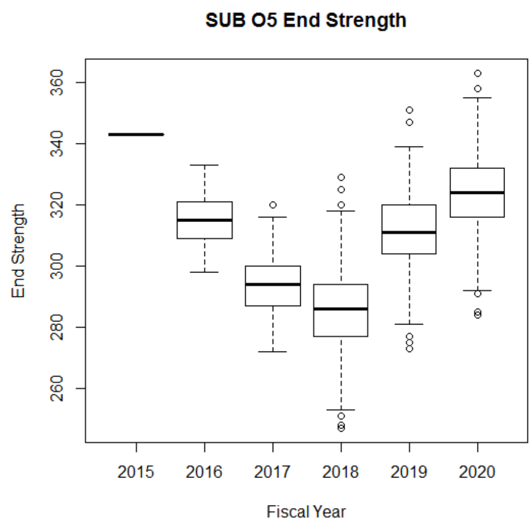
SUB O5 End Strength
End Strength
260
280
300
320
340
360
2015
2016
2017
2018
2019
2020
Fiscal Year

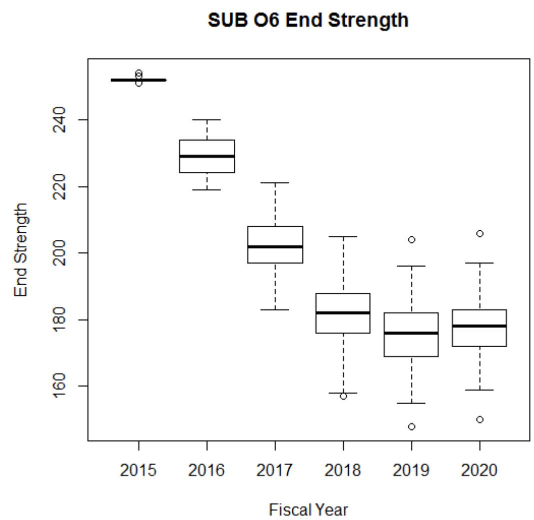
SUB O6 End Strength
End Strength
160
180
200
220
240
2015
2016
2017
2018
2019
2020
Fiscal Year

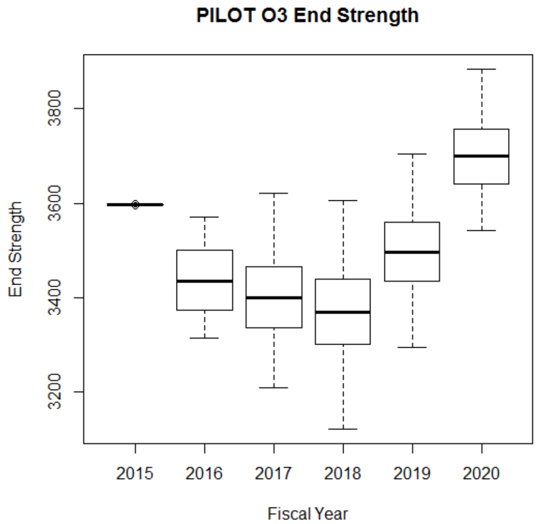
PILOT O3 End Strength
End Strength
3200
3400
3600
3800
2015
2016
2017
2018
2019
2020
Fiscal Year

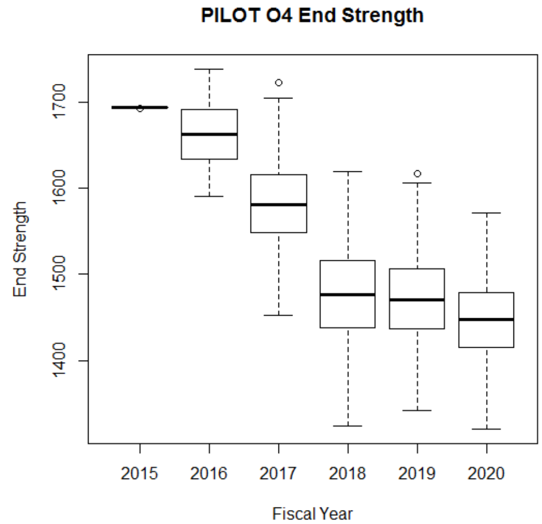
PILOT O4 End Strength
End Strength
1400
1500
1600
1700
2015
2016
2017
2018
2019
2020
Fiscal Year

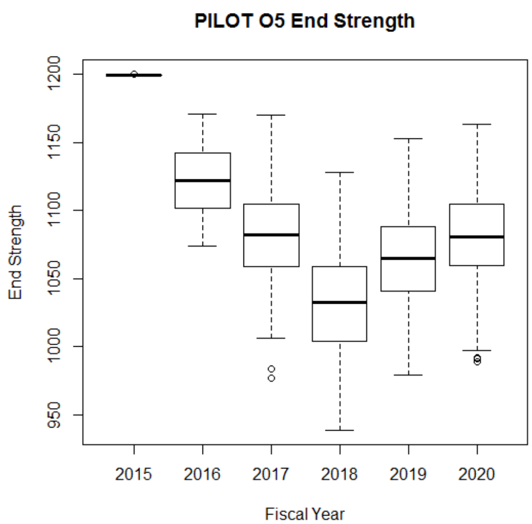
PILOT O5 End Strength
End Strength
950
1000
1050
1100
1150
1200
2015
2016
2017
2018
2019
2020
Fiscal Year

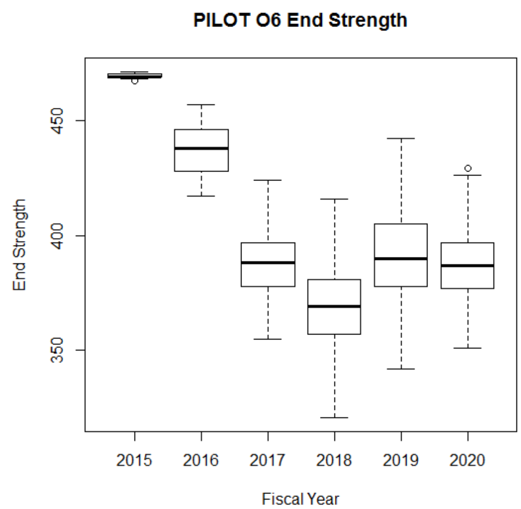
PILOT O6 End Strength
End Strength
350
400
450
2015
2016
2017
2018
2019
2020
Fiscal Year

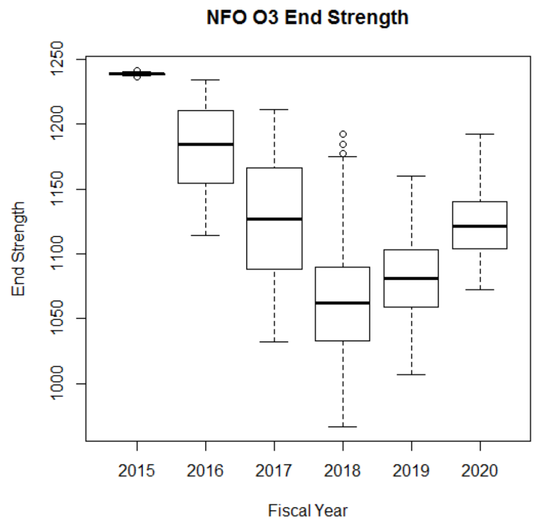
NFO O3 End Strength
End Strength
1000
1050
1100
1150
1200
1250
2015
2016
2017
2018
2019
2020
Fiscal Year

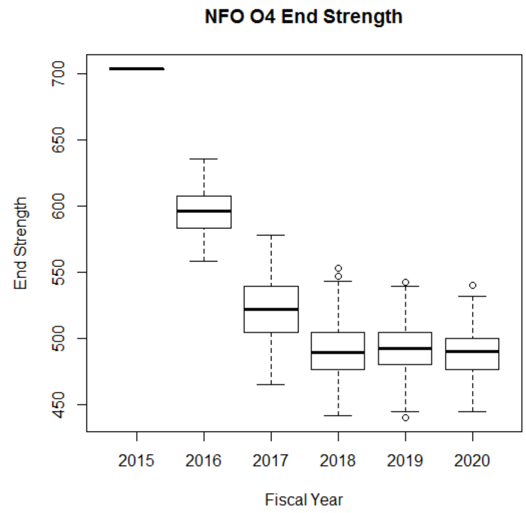
NFO O4 End Strength
End Strength
450
500
550
600
650
700
2015
2016
2017
2018
2019
2020
Fiscal Year

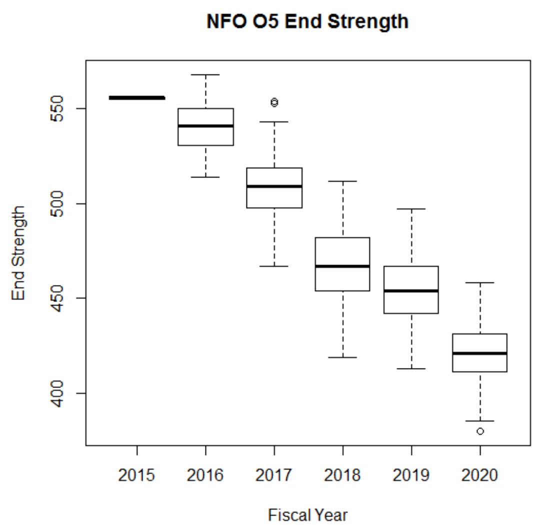
NFO O5 End Strength
End Strength
400
450
500
550
2015
2016
2017
2018
2019
2020
Fiscal Year

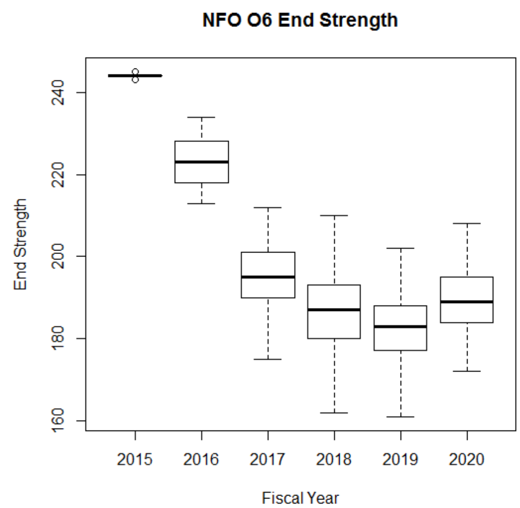
NFO O6 End Strength
End Strength
160
180
200
220
240
2015
2016
2017
2018
2019
2020
Fiscal Year

APPENDIX B. SUMMARIES OF REMAINING LINEAR REGRESSION MODELS

SWO O3

Summary of Fit

RSquare	0.963771
RSquare Adj	0.962451
Root Mean Square Error	8.885353
Mean of Response	2883.272
Observations (or Sum Wgts)	257

Analysis of Variance

Parameter Estimates

Term	Estimate	Std Error	t Ratio	Prob>|t|
Intercept	3180.9827	4.348211	731.56	<.0001*
SWO_O3_y1	-0.155018	0.009195	-16.86	<.0001*
SWO_O3_y2	-0.300314	0.009195	-32.66	<.0001*
SWO_O3_y3	-0.47573	0.009195	-51.74	<.0001*
SUB_O3_y1	-0.155183	0.019919	-7.79	<.0001*
SUB_O3_y2	-0.15308	0.01992	-7.68	<.0001*
SUB_O3_y3	-0.118925	0.019919	-5.97	<.0001*
AV_O3_y1	-0.154193	0.005283	-29.19	<.0001*
AV_O3_y2	-0.152374	0.005283	-28.84	<.0001*
AV_O3_y3	-0.140675	0.005283	-26.63	<.0001*

SWO O5

Summary of Fit

RSquare	0.969627
RSquare Adj	0.967467
Root Mean Square Error	5.539391
Mean of Response	743.1556
Observations (or Sum Wgts)	257

Analysis of Variance

Parameter Estimates

Term	Estimate	Std Error	t Ratio	Prob>|t|
Intercept	812.22615	3.705235	219.21	<.0001*
SWO_O3_y1	-0.04285	0.005732	-7.47	<.0001*
SWO_O4_y1	-0.267703	0.013397	-19.98	<.0001*
SWO_O4_y2	-0.312782	0.013397	-23.35	<.0001*
SWO_O4_y3	-0.292743	0.013397	-21.85	<.0001*
SWO_O5O6_y1	-0.3135	0.01157	-27.10	<.0001*
SWO_O5O6_y2	-0.335359	0.01157	-28.98	<.0001*
SWO_O5O6_y3	-0.411395	0.01157	-35.56	<.0001*
SUB_O4_y1	0.1870274	0.027769	6.74	<.0001*
SUB_O4_y2	0.2129429	0.027769	7.67	<.0001*
SUB_O4_y3	0.2427461	0.027769	8.74	<.0001*
AV_O3_y1	0.0192533	0.003294	5.85	<.0001*
AV_O4_y1	0.1821736	0.006513	27.97	<.0001*
AV_O4_y2	0.2158194	0.006513	33.14	<.0001*
AAV_O4_y3	0.1919986	0.006513	29.48	<.0001*
AV_O5O6_y1	-0.061598	0.006588	-9.35	<.0001*
AV_O5O6_y2	-0.075572	0.006588	-11.47	<.0001*
AV_O5O6_y3	-0.076976	0.006588	-11.68	<.0001*

SWO O6

Summary of Fit

RSquare	0.967562
RSquare Adj	0.966243
Root Mean Square Error	3.578905
Mean of Response	425.2023
Observations (or Sum Wgts)	257

Analysis of Variance

Parameter Estimates

Term	Estimate	Std Error	t Ratio	Prob>|t|
Intercept	475.51279	1.843743	257.91	<.0001*
SWO_O4_y1	-0.02374	0.008655	-2.74	0.0065*
SWO_O5O6_y1	-0.290992	0.007475	-38.93	<.0001*
SWO_O5O6_y2	-0.371285	0.007475	-49.67	<.0001*
SWO_O5O6_y3	-0.382221	0.007475	-51.13	<.0001*
SUB_O5O6_y1	0.0621541	0.017941	3.46	0.0006*
SUB_O5O6_y2	0.0615723	0.017941	3.43	0.0007*
SUB_O5O6_y3	0.0722164	0.017941	4.03	<.0001*
AV_O5O6_y1	0.054776	0.004256	12.87	<.0001*
AV_O5O6_y2	0.0702806	0.004256	16.51	<.0001*
AV_O5O6_y3	0.067137	0.004256	15.77	<.0001*

Submarine Officer O4

Summary of Fit

RSquare	0.977517
RSquare Adj	0.975817
Root Mean Square Error	4.058447
Mean of Response	528.2724
Observations (or Sum Wgts)	257

Analysis of Variance

Parameter Estimates

Term	Estimate	Std Error	t Ratio	Prob>|t|
Intercept	641.60409	2.796947	229.39	<.0001*
SWO_O3_y1	0.037499	0.0042	8.93	<.0001*
SWO_O3_y2	0.0460422	0.0042	10.96	<.0001*
SWO_O3_y3	0.0403426	0.0042	9.61	<.0001*
SWO_O4_y1	-0.060997	0.009815	-6.21	<.0001*
SWO_O4_y2	-0.064637	0.009815	-6.59	<.0001*
SWO_O4_y3	-0.048654	0.009815	-4.96	<.0001*
SUB_O3_y1	-0.541569	0.009099	-59.52	<.0001*
SUB_O3_y2	-0.439905	0.009099	-48.35	<.0001*
SUB_O3_y3	-0.341045	0.009099	-37.48	<.0001*
SUB_O4_y1	-0.327364	0.020345	-16.09	<.0001*
SUB_O4_y2	-0.357429	0.020345	-17.57	<.0001*
SUB_O4_y3	-0.476409	0.020345	-23.42	<.0001*
AV_O3_y1	0.0405425	0.002413	16.80	<.0001*
AV_O3_y2	0.0473556	0.002413	19.63	<.0001*
AV_O3_y3	0.0486883	0.002413	20.18	<.0001*
AV_O4_y1	-0.05286	0.004772	-11.08	<.0001*
AV_O4_y2	-0.066722	0.004772	-13.98	<.0001*
AAV_O4_y3	-0.064868	0.004772	-13.59	<.0001*

Submarine Officer O5

Summary of Fit

RSquare	0.944433
RSquare Adj	0.939467
Root Mean Square Error	3.314659
Mean of Response	323.642
Observations (or Sum Wgts)	257

Analysis of Variance

Parameter Estimates

Term	Estimate	Std Error	t Ratio	Prob>|t|
Intercept	373.83011	2.47098	151.29	<.0001*
SWO_O4_y1	0.0563208	0.008016	7.03	<.0001*
SWO_O4_y2	0.0722079	0.008016	9.01	<.0001*
SWO_O4_y3	0.0460574	0.008016	5.75	<.0001*
SWO_O5O6_y1	-0.031685	0.006924	-4.58	<.0001*
SWO_O5O6_y2	-0.024875	0.006924	-3.59	0.0004*
SWO_O5O6_y3	-0.029268	0.006924	-4.23	<.0001*
SUB_O3_y1	-0.046998	0.007431	-6.32	<.0001*
SUB_O4_y1	-0.351197	0.016617	-21.14	<.0001*
SUB_O4_y2	-0.450062	0.016617	-27.09	<.0001*
SUB_O4_y3	-0.406661	0.016617	-24.47	<.0001*
SUB_O5O6_y1	-0.264246	0.016617	-15.90	<.0001*
SUB_O5O6_y2	-0.297403	0.016617	-17.90	<.0001*
SUB_O5O6_y3	-0.383469	0.016616	-23.08	<.0001*
AV_O3_y1	0.0093713	0.001971	4.76	<.0001*
AV_O3_y2	0.0061295	0.001971	3.11	0.0021*
AV_O4_y1	0.0507659	0.003897	13.03	<.0001*
AV_O4_y2	0.0650879	0.003897	16.70	<.0001*
AAV_O4_y3	0.0618496	0.003897	15.87	<.0001*
AV_O5O6_y1	-0.021116	0.003942	-5.36	<.0001*
AV_O5O6_y2	-0.024282	0.003942	-6.16	<.0001*
AV_O5O6_y3	-0.030846	0.003942	-7.82	<.0001*

Submarine Officer O6

Summary of Fit

RSquare	0.936973
RSquare Adj	0.934676
Root Mean Square Error	2.2877
Mean of Response	177.4786
Observations (or Sum Wgts)	257

Analysis of Variance

Parameter Estimates

Term	Estimate	Std Error	t Ratio	Prob>|t|
Intercept	197.23558	1.119417	176.19	<.0001*
SWO_O5O6_y1	0.0247078	0.004778	5.17	<.0001*
SWO_O5O6_y2	0.0331119	0.004778	6.93	<.0001*
SWO_O5O6_y3	0.0313243	0.004778	6.56	<.0001*
SUB_O5O6_y1	-0.336251	0.011468	-29.32	<.0001*
SUB_O5O6_y2	-0.396954	0.011468	-34.61	<.0001*
SUB_O5O6_y3	-0.404559	0.011468	-35.28	<.0001*
AV_O5O6_y1	0.0209292	0.002721	7.69	<.0001*
AV_O5O6_y2	0.0283444	0.002721	10.42	<.0001*
AV_O5O6_y3	0.0274072	0.002721	10.07	<.0001*

Pilot O3

Summary of Fit

RSquare	0.990085
RSquare Adj	0.989724
Root Mean Square Error	7.750707
Mean of Response	3701.934
Observations (or Sum Wgts)	257

Analysis of Variance

Parameter Estimates

Term	Estimate	Std Error	t Ratio	Prob>|t|
Intercept	4150.7035	3.792951	1094.3	<.0001*
SWO_O3_y1	-0.22636	0.008021	-28.22	<.0001*
SWO_O3_y2	-0.198672	0.008021	-24.77	<.0001*
SWO_O3_y3	-0.150727	0.008021	-18.79	<.0001*
SUB_O3_y1	-0.20567	0.017376	-11.84	<.0001*
SUB_O3_y2	-0.163109	0.017376	-9.39	<.0001*
SUB_O3_y3	-0.125899	0.017376	-7.25	<.0001*
AV_O3_y1	-0.329099	0.004608	-71.42	<.0001*
AV_O3_y2	-0.387686	0.004608	-84.13	<.0001*
AV_O3_y3	-0.471749	0.004608	-102.4	<.0001*

Pilot O4

Summary of Fit

RSquare	0.971173
RSquare Adj	0.969251
Root Mean Square Error	8.889062
Mean of Response	1447.599
Observations (or Sum Wgts)	257

Analysis of Variance

Parameter Estimates

Term	Estimate	Std Error	t Ratio	Prob>|t|
Intercept	1629.3064	5.772588	282.25	<.0001*
SWO_O3_y1	0.2238367	0.009199	24.33	<.0001*
SWO_O3_y2	0.2093452	0.009199	22.76	<.0001*
SWO_O3_y3	0.1625023	0.009199	17.67	<.0001*
SWO_O4_y1	-0.140177	0.021497	-6.52	<.0001*
SWO_O4_y2	-0.183977	0.021497	-8.56	<.0001*
SWO_O4_y3	-0.158127	0.021497	-7.36	<.0001*
SUB_O3_y1	0.1995977	0.019928	10.02	<.0001*
SUB_O3_y2	0.1787451	0.019928	8.97	<.0001*
SUB_O3_y3	0.1401367	0.019928	7.03	<.0001*
SUB_O4_y2	-0.166014	0.04456	-3.73	0.0002*
AV_O3_y1	-0.148328	0.005285	-28.07	<.0001*
AV_O3_y2	-0.192946	0.005285	-36.51	<.0001*
AV_O3_y3	-0.164323	0.005285	-31.09	<.0001*
AV_O4_y1	-0.278814	0.010451	-26.68	<.0001*
AV_O4_y2	-0.312803	0.010451	-29.93	<.0001*
AAV_O4_y3	-0.412197	0.010451	-39.44	<.0001*

Pilot O6

Summary of Fit

RSquare	0.953161
RSquare Adj	0.951454
Root Mean Square Error	3.219788
Mean of Response	387.1868
Observations (or Sum Wgts)	257

Analysis of Variance

Parameter Estimates

Term	Estimate	Std Error	t Ratio	Prob>|t|
Intercept	419.87026	1.575507	266.50	<.0001*
SWO_O5O6_y1	0.0758685	0.006725	11.28	<.0001*
SWO_O5O6_y2	0.0923955	0.006725	13.74	<.0001*
SWO_O5O6_y3	0.0875866	0.006725	13.02	<.0001*
SUB_O5O6_y1	0.0471023	0.01614	2.92	0.0038*
SUB_O5O6_y2	0.0960962	0.016141	5.95	<.0001*
SUB_O5O6_y3	0.0585211	0.016141	3.63	0.0003*
AV_O5O6_y1	-0.114842	0.003829	-29.99	<.0001*
AV_O5O6_y2	-0.153611	0.003829	-40.12	<.0001*
AV_O5O6_y3	-0.170713	0.003829	-44.58	<.0001*

NFO O3

Summary of Fit

RSquare	0.975853
RSquare Adj	0.974973
Root Mean Square Error	3.90896
Mean of Response	1122.755
Observations (or Sum Wgts)	257

Analysis of Variance

Parameter Estimates

Term	Estimate	Std Error	t Ratio	Prob>|t|
Intercept	1264.8452	1.912922	661.21	<.0001*
SWO_O3_y1	-0.063615	0.004045	-15.73	<.0001*
SWO_O3_y2	-0.060232	0.004045	-14.89	<.0001*
SWO_O3_y3	-0.043142	0.004045	-10.67	<.0001*
SUB_O3_y1	-0.059903	0.008763	-6.84	<.0001*
SUB_O3_y2	-0.057253	0.008763	-6.53	<.0001*
SUB_O3_y3	-0.037967	0.008763	-4.33	<.0001*
AV_O3_y1	-0.112428	0.002324	-48.38	<.0001*
AV_O3_y2	-0.123538	0.002324	-53.16	<.0001*
AV_O3_y3	-0.14953	0.002324	-64.34	<.0001*

NFO O4

Summary of Fit

RSquare	0.947523
RSquare Adj	0.944257
Root Mean Square Error	4.322463
Mean of Response	488.3346
Observations (or Sum Wgts)	257

Analysis of Variance

Parameter Estimates

Term	Estimate	Std Error	t Ratio	Prob>|t|
Intercept	559.87221	2.716365	206.11	<.0001*
SWO_O3_y1	0.0612439	0.004473	13.69	<.0001*
SWO_O3_y2	0.063731	0.004473	14.25	<.0001*
SWO_O3_y3	0.053948	0.004473	12.06	<.0001*
SWO_O4_y1	-0.043382	0.010453	-4.15	<.0001*
SWO_O4_y2	-0.047955	0.010453	-4.59	<.0001*
SWO_O4_y3	-0.032555	0.010453	-3.11	0.0021*
SUB_O3_y1	0.0573343	0.00969	5.92	<.0001*
SUB_O3_y2	0.0587422	0.00969	6.06	<.0001*
SUB_O3_y3	0.042399	0.00969	4.38	<.0001*
AV_O3_y1	-0.059736	0.00257	-23.24	<.0001*
AV_O3_y2	-0.070119	0.00257	-27.28	<.0001*
AV_O3_y3	-0.072911	0.00257	-28.37	<.0001*
AV_O4_y1	-0.097717	0.005082	-19.23	<.0001*
AV_O4_y2	-0.116785	0.005082	-22.98	<.0001*
AAV_O4_y3	-0.133355	0.005082	-26.24	<.0001*

NFO O5

Summary of Fit

RSquare	0.942832
RSquare Adj	0.939273
Root Mean Square Error	3.638338
Mean of Response	420.5875
Observations (or Sum Wgts)	257

Analysis of Variance

Parameter Estimates

Term	Estimate	Std Error	t Ratio	Prob>|t|
Intercept	502.18706	2.284713	219.80	<.0001*
SWO_O3_y1	0.0143452	0.003765	3.81	0.0002*
SWO_O4_y1	0.0535224	0.008799	6.08	<.0001*
SWO_O4_y2	0.054896	0.008799	6.24	<.0001*
SWO_O4_y3	0.0551468	0.008799	6.27	<.0001*
SWO_O5O6_y1	-0.032262	0.0076	-4.25	<.0001*
SWO_O5O6_y2	-0.039374	0.0076	-5.18	<.0001*
SWO_O5O6_y3	-0.033574	0.0076	-4.42	<.0001*
SUB_O4_y1	0.0480874	0.018239	2.64	0.0089*
AV_O3_y2	-0.005757	0.002163	-2.66	0.0083*
AV_O4_y1	-0.097517	0.004278	-22.80	<.0001*
AV_O4_y2	-0.097194	0.004278	-22.72	<.0001*
AAV_O4_y3	-0.085859	0.004278	-20.07	<.0001*
AV_O5O6_y1	-0.101831	0.004327	-23.53	<.0001*
AV_O5O6_y2	-0.114248	0.004327	-26.40	<.0001*
AV_O5O6_y3	-0.142727	0.004327	-32.99	<.0001*

LIST OF REFERENCES

Borozny, E.E. (2015). Projecting Navy officer inventory with data farming (Master's thesis). Unpublished manuscript, Operations Research Department, Naval Postgraduate School, Monterey, CA.

Cioppa, T. M., & Lucas, T. W. (2007). Efficient nearly orthogonal and space-filling Latin hypercubes. *Technometrics*, *49*(1), 45–55.

Clark, D. T. (2009). Navy officer manpower optimization incorporating budgetary constraints (Master's thesis). Retrieved from Calhoun https://calhoun.nps.edu/handle/10945/4892

Crawley, M. J. (2013). *The R Book* (2nd ed.). West Sussex, UK: John Wiley and Sons.

Department of the Navy. (2015, Jun. 9). Department of the Navy end strength management (SECNAV Instruction 1080.4). Washington, DC: Author.

Department of the Navy. (2013, Mar. 26). Highlights of the Department of the Navy FY 2014 Budget. Washington, DC: Author.

Department of the Navy. (2013, November). Manpower, personnel, training and education modeling catalog. Washington, DC: Author. Retrieved from https://dodlearn.nps.edu/access/content/group/366940ea-44b7-43fb-9b96-c271a2ecd54a/HR%20Community%20Documents/MPTE%20Models%20Catalog%20-%2015%20Nov%202013.pdf

Defense Officer Personnel Management Act of 1980, Pub. L. 96–513, 94 Stat. 2835 (1980). Retrieved from http://www.gpo.gov/fdsys/pkg/STATUTE-94/pdf/STATUTE-94-Pg2835.pdf

Erdman, R. W. (2010). Using experimental design and data analysis to study the enlisted specialty model for the U.S. Army G1 (Master's thesis). Retrieved from Calhoun https://calhoun.nps.edu/handle/10945/5313

Faraway, J. J. (2005). *Linear models with R*. Boca Raton, FL: CRC Press.

Horne, G., & Meyer, T. (2010). *Data farming and defense applications*. Paper presented at MODSIM World Conference and Expo, Hampton, VA. Retrieved from https://harvest.nps.edu/papers/MODSIMPaper-HorneMeyer.pdf

JMP Pro, Version 10.0.0 [Computer software] (2013). Cary, NC: SAS Institute, Inc.

Kleijnen, J. P. C., Sanchez, S. M., Lucas, T. W., & Cioppa, T. M. (2005). A user's guide to the brave new world of designing simulation experiments. *INFORMS Journal on Computing*, *17*(3), 263–289.

Lucas, T. W. (2000). Stochastic versus deterministic argument for combat simulations: Tales of when the average won't do! *Military Operations Research: Special Issue on Warfare Analysis and Complexity-the New Sciences*, *5*(3), 9–28.

Moran, W. (2014, December). Once again…a moment ripe for change. *Proceedings* magazine. Retrieved from http://www.usni.org/magazines/proceedings/2014-12/once-again-moment-ripe-change

Mundy, J. (2014). *OSAM Documentation*. Annadale, VA: SAG Corporation.

R Core Team (2015). R: A language and environment for statistical computing [Computer software]. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from http://www.R-project.org

Sanchez, S. M. (2011). NOLHdesigns [Spreadsheet]. Retrieved from https://harvest.nps.edu/software.html

Sanchez, S. M., & Wan, H. (2012). Work smarter, not harder: Tutorial on designing and conducting simulation experiments. *In Proceedings of the 2012 Winter Simulation Conference*. December 9–12, Berlin, Germany.

Schank, J. F., Harrell, M. C., & Thie, H. J. (1997). Relating resources to personnel readiness: Use of Army strength management models. Santa Monica, CA: RAND.

Sibley, C. N. (2012). Analyzing Navy officer inventory projection using data farming (Master's thesis). Retrieved from https://calhoun.nps.edu/handle/10945/6868

Snodgrass, G., & Kohlmann, B. (2014). *2014 Navy retention study*. Washington, DC: U.S. Navy. Retrieved from http://www.dodretention.org/

Venables, W. N., & Ripley, B. D. (2002). *Modern applied statistics with S* (4th ed.). New York, NY: Springer.

Weiss, H., MacDermid, S., Strauss, R., Kurek, K., Le, B., & Robbins, D. (2002). Retention in the armed forces: past approaches and new research directions. West Lafayette, IN: Military Family Research Institute. Retrieved from https://www.mfri.purdue.edu/resources/public/reports/Retention%20in%20the%20Armed%20Forces.pdf

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center
   Ft. Belvoir, Virginia

2. Dudley Knox Library
   Naval Postgraduate School
   Monterey, California